



Catalyst 2970 Switch Software Configuration Guide

Cisco IOS Release 12.1(19)EA1
October 2003

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: DOC-7815462=
Text Part Number: 78-15462-03



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, CCSP, the Cisco Arrow logo, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0304R)

Catalyst 2970 Switch Software Configuration Guide
Copyright ©2003 Cisco Systems, Inc. All rights reserved.



Preface xxvii

Audience xxvii

Purpose xxvii

Conventions xxviii

Related Publications xxix

Obtaining Documentation xxix

Cisco.com xxix

Documentation CD-ROM xxix

Ordering Documentation xxx

Documentation Feedback xxx

Obtaining Technical Assistance xxx

Cisco TAC Website xxx

Opening a TAC Case xxxi

TAC Case Priority Definitions xxxi

Obtaining Additional Publications and Information xxxi

CHAPTER 1

Overview 1-1

Features 1-1

Default Settings After Initial Switch Configuration 1-8

Network Configuration Examples 1-10

Design Concepts for Using the Switch 1-10

Small to Medium-Sized Network Using Catalyst 2970 Switches 1-14

Long-Distance, High-Bandwidth Transport Configuration 1-15

Where to Go Next 1-15

CHAPTER 2

Using the Command-Line Interface 2-1

Understanding Command Modes 2-1

Understanding the Help System 2-3

Understanding Abbreviated Commands 2-4

Understanding no and default Forms of Commands 2-4

Understanding CLI Error Messages 2-5

Using Command History	2-5
Changing the Command History Buffer Size	2-5
Recalling Commands	2-6
Disabling the Command History Feature	2-6
Using Editing Features	2-6
Enabling and Disabling Editing Features	2-7
Editing Commands through Keystrokes	2-7
Editing Command Lines that Wrap	2-8
Searching and Filtering Output of show and more Commands	2-9
Accessing the CLI	2-10
Accessing the CLI through a Console Connection or through Telnet	2-10
Accessing the CLI from a Browser	2-10

CHAPTER 3

Getting Started with CMS	3-1
Understanding CMS	3-1
Front Panel View	3-2
Topology View	3-2
CMS Menu Bar, Toolbar, and Feature Bar	3-2
Online Help	3-4
Configuration Modes	3-5
Guide Mode	3-5
Expert Mode	3-6
Wizards	3-6
Privilege Levels	3-6
Access to Older Switches In a Cluster	3-7
Configuring CMS	3-7
CMS Requirements	3-7
Minimum Hardware Configuration	3-7
Operating System and Browser Support	3-8
Browser Plug-In Requirements	3-8
Cross-Platform Considerations	3-9
HTTP Access to CMS	3-9
Specifying an HTTP Port (Nondefault Configuration Only)	3-9
Configuring an Authentication Method (Nondefault Configuration Only)	3-10
Displaying CMS	3-10
Launching CMS	3-10
Front Panel View	3-13

Topology View	3-14
CMS Icons	3-15
Where to Go Next	3-15

CHAPTER 4

Assigning the Switch IP Address and Default Gateway 4-1

Understanding the Boot Process	4-1
Assigning Switch Information	4-2
Default Switch Information	4-3
Understanding DHCP-Based Autoconfiguration	4-3
DHCP Client Request Process	4-4
Configuring DHCP-Based Autoconfiguration	4-5
Configuring the DHCP Server	4-5
Configuring the TFTP Server	4-6
Configuring the DNS	4-6
Configuring the Relay Device	4-6
Obtaining Configuration Files	4-7
Example Configuration	4-8
Manually Assigning IP Information	4-10
Checking and Saving the Running Configuration	4-11
Modifying the Startup Configuration	4-12
Default Boot Configuration	4-12
Automatically Downloading a Configuration File	4-13
Specifying the Filename to Read and Write the System Configuration	4-13
Booting Manually	4-13
Booting a Specific Software Image	4-14
Controlling Environment Variables	4-15
Scheduling a Reload of the Software Image	4-17
Configuring a Scheduled Reload	4-17
Displaying Scheduled Reload Information	4-18

CHAPTER 5

Clustering Switches 5-1

Understanding Switch Clusters	5-2
Cluster Command Switch Characteristics	5-3
Standby Cluster Command Switch Characteristics	5-3
Candidate Switch and Cluster Member Switch Characteristics	5-4
Planning a Switch Cluster	5-4
Automatic Discovery of Cluster Candidates and Members	5-5
Discovery Through CDP Hops	5-5
Discovery Through Non-CDP-Capable and Noncluster-Capable Devices	5-6

Discovery Through Different VLANs	5-7
Discovery Through Different Management VLANs	5-8
Discovery of Newly Installed Switches	5-9
HSRP and Standby Cluster Command Switches	5-10
Virtual IP Addresses	5-11
Other Considerations for Cluster Standby Groups	5-11
Automatic Recovery of Cluster Configuration	5-12
IP Addresses	5-13
Host Names	5-13
Passwords	5-14
SNMP Community Strings	5-14
TACACS+ and RADIUS	5-14
Access Modes in CMS	5-15
LRE Profiles	5-15
Availability of Switch-Specific Features in Switch Clusters	5-15
Creating a Switch Cluster	5-15
Enabling a Cluster Command Switch	5-16
Adding Cluster Member Switches	5-16
Creating a Cluster Standby Group	5-18
Verifying a Switch Cluster	5-20
Using the CLI to Manage Switch Clusters	5-21
Catalyst 1900 and Catalyst 2820 CLI Considerations	5-21
Using SNMP to Manage Switch Clusters	5-21

CHAPTER 6

Administering the Switch 6-1

Managing the System Time and Date	6-1
Understanding the System Clock	6-2
Understanding Network Time Protocol	6-2
Configuring NTP	6-4
Default NTP Configuration	6-4
Configuring NTP Authentication	6-5
Configuring NTP Associations	6-6
Configuring NTP Broadcast Service	6-7
Configuring NTP Access Restrictions	6-8
Configuring the Source IP Address for NTP Packets	6-10
Displaying the NTP Configuration	6-11
Configuring Time and Date Manually	6-11
Setting the System Clock	6-12
Displaying the Time and Date Configuration	6-12

Configuring the Time Zone	6-13
Configuring Summer Time (Daylight Saving Time)	6-14
Configuring a System Name and Prompt	6-16
Default System Name and Prompt Configuration	6-16
Configuring a System Name	6-16
Configuring a System Prompt	6-17
Understanding DNS	6-17
Default DNS Configuration	6-18
Setting Up DNS	6-18
Displaying the DNS Configuration	6-19
Creating a Banner	6-19
Default Banner Configuration	6-19
Configuring a Message-of-the-Day Login Banner	6-20
Configuring a Login Banner	6-21
Managing the MAC Address Table	6-22
Building the Address Table	6-22
MAC Addresses and VLANs	6-23
Default MAC Address Table Configuration	6-23
Changing the Address Aging Time	6-23
Removing Dynamic Address Entries	6-24
Configuring MAC Address Notification Traps	6-24
Adding and Removing Static Address Entries	6-26
Configuring Unicast MAC Address Filtering	6-27
Displaying Address Table Entries	6-28
Managing the ARP Table	6-29

CHAPTER 7

Configuring Switch-Based Authentication	7-1
Preventing Unauthorized Access to Your Switch	7-1
Protecting Access to Privileged EXEC Commands	7-2
Default Password and Privilege Level Configuration	7-2
Setting or Changing a Static Enable Password	7-3
Protecting Enable and Enable Secret Passwords with Encryption	7-4
Disabling Password Recovery	7-5
Setting a Telnet Password for a Terminal Line	7-6
Configuring Username and Password Pairs	7-7
Configuring Multiple Privilege Levels	7-8
Setting the Privilege Level for a Command	7-8
Changing the Default Privilege Level for Lines	7-9
Logging into and Exiting a Privilege Level	7-10

Controlling Switch Access with TACACS+	7-10
Understanding TACACS+	7-10
TACACS+ Operation	7-12
Configuring TACACS+	7-13
Default TACACS+ Configuration	7-13
Identifying the TACACS+ Server Host and Setting the Authentication Key	7-13
Configuring TACACS+ Login Authentication	7-14
Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services	7-16
Starting TACACS+ Accounting	7-17
Displaying the TACACS+ Configuration	7-17
Controlling Switch Access with RADIUS	7-18
Understanding RADIUS	7-18
RADIUS Operation	7-19
Configuring RADIUS	7-20
Default RADIUS Configuration	7-20
Identifying the RADIUS Server Host	7-21
Configuring RADIUS Login Authentication	7-23
Defining AAA Server Groups	7-25
Configuring RADIUS Authorization for User Privileged Access and Network Services	7-27
Starting RADIUS Accounting	7-28
Configuring Settings for All RADIUS Servers	7-29
Configuring the Switch to Use Vendor-Specific RADIUS Attributes	7-29
Configuring the Switch for Vendor-Proprietary RADIUS Server Communication	7-31
Displaying the RADIUS Configuration	7-31
Controlling Switch Access with Kerberos	7-32
Understanding Kerberos	7-32
Kerberos Operation	7-34
Authenticating to a Boundary Switch	7-35
Obtaining a TGT from a KDC	7-35
Authenticating to Network Services	7-35
Configuring Kerberos	7-36
Configuring the Switch for Local Authentication and Authorization	7-36
Configuring the Switch for Secure Shell	7-37
Understanding SSH	7-38
SSH Servers, Integrated Clients, and Supported Versions	7-38
Limitations	7-38

Configuring SSH	7-39
Configuration Guidelines	7-39
Setting Up the Switch to Run SSH	7-39
Configuring the SSH Server	7-40
Displaying the SSH Configuration and Status	7-41

CHAPTER 8

Configuring 802.1X Port-Based Authentication 8-1

Understanding 802.1X Port-Based Authentication	8-1
Device Roles	8-2
Authentication Initiation and Message Exchange	8-3
Ports in Authorized and Unauthorized States	8-4
Supported Topologies	8-5
Using 802.1X with Port Security	8-6
Using 802.1X with Voice VLAN Ports	8-6
Using 802.1X with VLAN Assignment	8-7
Using 802.1X with Guest VLAN	8-8
Using 802.1X with Per-User ACLs	8-9
Configuring 802.1X Authentication	8-9
Default 802.1X Configuration	8-10
802.1X Configuration Guidelines	8-11
Upgrading from a Previous Software Release	8-12
Configuring 802.1X Authentication	8-12
Configuring the Switch-to-RADIUS-Server Communication	8-13
Configuring Periodic Re-Authentication	8-14
Manually Re-Authenticating a Client Connected to a Port	8-15
Changing the Quiet Period	8-15
Changing the Switch-to-Client Retransmission Time	8-16
Setting the Switch-to-Client Frame-Retransmission Number	8-17
Configuring the Host Mode	8-17
Configuring a Guest VLAN	8-18
Resetting the 802.1X Configuration to the Default Values	8-19
Displaying 802.1X Statistics and Status	8-19

CHAPTER 9

Configuring Interface Characteristics 9-1

Understanding Interface Types	9-1
Port-Based VLANs	9-2
Switch Ports	9-2
Access Ports	9-2
Trunk Ports	9-3

EtherChannel Port Groups	9-4
Connecting Interfaces	9-4
Using Interface Configuration Mode	9-5
Procedures for Configuring Interfaces	9-5
Configuring a Range of Interfaces	9-6
Configuring and Using Interface Range Macros	9-7
Configuring Ethernet Interfaces	9-9
Default Ethernet Interface Configuration	9-9
Configuring Interface Speed and Duplex Mode	9-10
Configuration Guidelines	9-10
Setting the Interface Speed and Duplex Parameters	9-11
Configuring IEEE 802.3X Flow Control	9-12
Configuring Auto-MDIX on an Interface	9-13
Adding a Description for an Interface	9-14
Configuring the System MTU	9-14
Monitoring and Maintaining the Interfaces	9-16
Monitoring Interface Status	9-16
Clearing and Resetting Interfaces and Counters	9-17
Shutting Down and Restarting the Interface	9-17

CHAPTER 10

Configuring SmartPort Macros	10-1
Understanding SmartPort Macros	10-1
Configuring Smart-Port Macros	10-2
Default SmartPort Macro Configuration	10-2
SmartPort Macro Configuration Guidelines	10-2
Creating and Applying SmartPort Macros	10-3
Displaying SmartPort Macros	10-4

CHAPTER 11

Configuring VLANs	11-1
Understanding VLANs	11-1
Supported VLANs	11-3
VLAN Port Membership Modes	11-3
Configuring Normal-Range VLANs	11-4
Token Ring VLANs	11-5
Normal-Range VLAN Configuration Guidelines	11-6
VLAN Configuration Mode Options	11-6
VLAN Configuration in config-vlan Mode	11-6
VLAN Configuration in VLAN Database Configuration Mode	11-7

Saving VLAN Configuration	11-7
Default Ethernet VLAN Configuration	11-7
Creating or Modifying an Ethernet VLAN	11-8
Deleting a VLAN	11-10
Assigning Static-Access Ports to a VLAN	11-11
Configuring Extended-Range VLANs	11-12
Default VLAN Configuration	11-12
Extended-Range VLAN Configuration Guidelines	11-12
Creating an Extended-Range VLAN	11-13
Displaying VLANs	11-14
Configuring VLAN Trunks	11-15
Trunking Overview	11-15
Encapsulation Types	11-16
802.1Q Configuration Considerations	11-17
Default Layer 2 Ethernet Interface VLAN Configuration	11-17
Configuring an Ethernet Interface as a Trunk Port	11-18
Interaction with Other Features	11-18
Configuring a Trunk Port	11-19
Defining the Allowed VLANs on a Trunk	11-20
Changing the Pruning-Eligible List	11-21
Configuring the Native VLAN for Untagged Traffic	11-21
Configuring Trunk Ports for Load Sharing	11-22
Load Sharing Using STP Port Priorities	11-22
Load Sharing Using STP Path Cost	11-24
Configuring VMPS	11-26
Understanding VMPS	11-26
Dynamic-Access Port VLAN Membership	11-27
Default VMPS Client Configuration	11-27
VMPS Configuration Guidelines	11-27
Configuring the VMPS Client	11-28
Entering the IP Address of the VMPS	11-28
Configuring Dynamic-Access Ports on VMPS Clients	11-29
Reconfirming VLAN Memberships	11-29
Changing the Reconfirmation Interval	11-29
Changing the Retry Count	11-30
Monitoring the VMPS	11-30
Troubleshooting Dynamic-Access Port VLAN Membership	11-31
VMPS Configuration Example	11-31

CHAPTER 12

Configuring VTP 12-1

Understanding VTP 12-1

The VTP Domain 12-2

VTP Modes 12-3

VTP Advertisements 12-3

VTP Version 2 12-4

VTP Pruning 12-4

Configuring VTP 12-6

Default VTP Configuration 12-6

VTP Configuration Options 12-6

VTP Configuration in Global Configuration Mode 12-7

VTP Configuration in VLAN Database Configuration Mode 12-7

VTP Configuration Guidelines 12-7

Domain Names 12-7

Passwords 12-8

VTP Version 12-8

Configuration Requirements 12-8

Configuring a VTP Server 12-9

Configuring a VTP Client 12-10

Disabling VTP (VTP Transparent Mode) 12-11

Enabling VTP Version 2 12-12

Enabling VTP Pruning 12-13

Adding a VTP Client Switch to a VTP Domain 12-14

Monitoring VTP 12-15

CHAPTER 13

Configuring Voice VLAN 13-1

Understanding Voice VLAN 13-1

Cisco IP Phone Voice Traffic 13-2

Cisco IP Phone Data Traffic 13-2

Configuring Voice VLAN 13-3

Default Voice VLAN Configuration 13-3

Voice VLAN Configuration Guidelines 13-3

Configuring a Port Connected to a Cisco 7960 IP Phone 13-4

Configuring IP Phone Voice Traffic 13-4

Configuring the Priority of Incoming Data Frames 13-5

Displaying Voice VLAN 13-6

CHAPTER 14

Configuring STP 14-1

Understanding Spanning-Tree Features	14-1
STP Overview	14-2
Spanning-Tree Topology and BPDUs	14-3
Bridge ID, Switch Priority, and Extended System ID	14-4
Spanning-Tree Interface States	14-4
Blocking State	14-5
Listening State	14-6
Learning State	14-6
Forwarding State	14-6
Disabled State	14-7
How a Switch or Port Becomes the Root Switch or Root Port	14-7
Spanning Tree and Redundant Connectivity	14-8
Spanning-Tree Address Management	14-8
Accelerated Aging to Retain Connectivity	14-8
Spanning-Tree Modes and Protocols	14-9
Supported Spanning-Tree Instances	14-9
Spanning-Tree Interoperability and Backward Compatibility	14-10
STP and IEEE 802.1Q Trunks	14-10
Configuring Spanning-Tree Features	14-11
Default Spanning-Tree Configuration	14-11
Spanning-Tree Configuration Guidelines	14-12
Changing the Spanning-Tree Mode	14-13
Disabling Spanning Tree	14-14
Configuring the Root Switch	14-14
Configuring a Secondary Root Switch	14-16
Configuring Port Priority	14-17
Configuring Path Cost	14-18
Configuring the Switch Priority of a VLAN	14-19
Configuring Spanning-Tree Timers	14-20
Configuring the Hello Time	14-20
Configuring the Forwarding-Delay Time for a VLAN	14-21
Configuring the Maximum-Aging Time for a VLAN	14-21
Displaying the Spanning-Tree Status	14-22

CHAPTER 15

Configuring MSTP 15-1

Understanding MSTP	15-2
Multiple Spanning-Tree Regions	15-2
IST, CIST, and CST	15-3

Operations Within an MST Region	15-3
Operations Between MST Regions	15-4
Hop Count	15-5
Boundary Ports	15-5
Interoperability with 802.1D STP	15-5
Understanding RSTP	15-6
Port Roles and the Active Topology	15-6
Rapid Convergence	15-7
Synchronization of Port Roles	15-8
Bridge Protocol Data Unit Format and Processing	15-9
Processing Superior BPDU Information	15-10
Processing Inferior BPDU Information	15-10
Topology Changes	15-10
Configuring MSTP Features	15-11
Default MSTP Configuration	15-12
MSTP Configuration Guidelines	15-12
Specifying the MST Region Configuration and Enabling MSTP	15-13
Configuring the Root Switch	15-14
Configuring a Secondary Root Switch	15-16
Configuring Port Priority	15-17
Configuring Path Cost	15-18
Configuring the Switch Priority	15-19
Configuring the Hello Time	15-19
Configuring the Forwarding-Delay Time	15-20
Configuring the Maximum-Aging Time	15-21
Configuring the Maximum-Hop Count	15-21
Specifying the Link Type to Ensure Rapid Transitions	15-22
Restarting the Protocol Migration Process	15-22
Displaying the MST Configuration and Status	15-23

CHAPTER 16

Configuring Optional Spanning-Tree Features 16-1

Understanding Optional Spanning-Tree Features	16-1
Understanding Port Fast	16-2
Understanding BPDU Guard	16-3
Understanding BPDU Filtering	16-3
Understanding UplinkFast	16-4
Understanding BackboneFast	16-5
Understanding Root Guard	16-7
Understanding Loop Guard	16-8

Configuring Optional Spanning-Tree Features	16-9
Default Optional Spanning-Tree Configuration	16-9
Optional Spanning-Tree Configuration Guidelines	16-9
Enabling Port Fast	16-10
Enabling BPDU Guard	16-11
Enabling BPDU Filtering	16-12
Enabling UplinkFast for Use with Redundant Links	16-13
Enabling BackboneFast	16-13
Enabling Root Guard	16-14
Enabling Loop Guard	16-15
Displaying the Spanning-Tree Status	16-16

CHAPTER 17

Configuring DHCP Features	17-1
Understanding DHCP Features	17-1
DHCP Snooping	17-1
Option-82 Data Insertion	17-2
Configuring DHCP Features	17-3
Default DHCP Configuration	17-3
DHCP Snooping Configuration Guidelines	17-3
Enabling DHCP Snooping and Option 82	17-4
Displaying DHCP Information	17-5
Displaying a Binding Table	17-5
Displaying the DHCP Snooping Configuration	17-5

CHAPTER 18

Configuring IGMP Snooping and MVR	18-1
Understanding IGMP Snooping	18-2
IGMP Versions	18-3
Joining a Multicast Group	18-3
Leaving a Multicast Group	18-5
Immediate-Leave Processing	18-6
IGMP Report Suppression	18-6
Configuring IGMP Snooping	18-6
Default IGMP Snooping Configuration	18-7
Enabling or Disabling IGMP Snooping	18-7
Setting the Snooping Method	18-8
Configuring a Multicast Router Port	18-9
Configuring a Host Statically to Join a Group	18-10
Enabling IGMP Immediate-Leave Processing	18-10
Disabling IGMP Report Suppression	18-11

Displaying IGMP Snooping Information	18-12
Understanding Multicast VLAN Registration	18-13
Using MVR in a Multicast Television Application	18-14
Configuring MVR	18-15
Default MVR Configuration	18-16
MVR Configuration Guidelines and Limitations	18-16
Configuring MVR Global Parameters	18-16
Configuring MVR Interfaces	18-18
Displaying MVR Information	18-19
Configuring IGMP Filtering and Throttling	18-20
Default IGMP Filtering and Throttling Configuration	18-21
Configuring IGMP Profiles	18-21
Applying IGMP Profiles	18-22
Setting the Maximum Number of IGMP Groups	18-23
Configuring the IGMP Throttling Action	18-24
Displaying IGMP Filtering and Throttling Configuration	18-25

CHAPTER 19

Configuring Port-Based Traffic Control	19-1
Configuring Storm Control	19-1
Understanding Storm Control	19-2
Default Storm Control Configuration	19-3
Enabling Storm Control	19-3
Configuring Protected Ports	19-5
Default Protected Port Configuration	19-5
Protected Port Configuration Guidelines	19-5
Configuring a Protected Port	19-5
Configuring Port Blocking	19-6
Default Port Blocking Configuration	19-6
Blocking Flooded Traffic on an Interface	19-6
Configuring Port Security	19-7
Understanding Port Security	19-7
Secure MAC Addresses	19-7
Security Violations	19-8
Default Port Security Configuration	19-9
Configuration Guidelines	19-9
Enabling and Configuring Port Security	19-10
Enabling and Configuring Port Security Aging	19-13
Displaying Port-Based Traffic Control Settings	19-15

CHAPTER 20

Configuring CDP 20-1

Understanding CDP 20-1

Configuring CDP 20-2

Default CDP Configuration 20-2

Configuring the CDP Characteristics 20-2

Disabling and Enabling CDP 20-3

Disabling and Enabling CDP on an Interface 20-4

Monitoring and Maintaining CDP 20-5

CHAPTER 21

Configuring UDLD 21-1

Understanding UDLD 21-1

Modes of Operation 21-1

Methods to Detect Unidirectional Links 21-2

Configuring UDLD 21-4

Default UDLD Configuration 21-4

Configuration Guidelines 21-4

Enabling UDLD Globally 21-5

Enabling UDLD on an Interface 21-5

Resetting an Interface Disabled by UDLD 21-6

Displaying UDLD Status 21-6

CHAPTER 22

Configuring SPAN and RSPAN 22-1

Understanding SPAN and RSPAN 22-1

Local SPAN 22-2

Remote SPAN 22-2

SPAN and RSPAN Concepts and Terminology 22-3

SPAN Sessions 22-3

Monitored Traffic 22-4

Source Ports 22-5

Source VLANs 22-6

VLAN Filtering 22-6

Destination Port 22-7

RSPAN VLAN 22-8

SPAN and RSPAN Interaction with Other Features 22-8

Configuring SPAN and RSPAN 22-9

Default SPAN and RSPAN Configuration 22-9

Configuring Local SPAN 22-10

SPAN Configuration Guidelines 22-10

Creating a Local SPAN Session	22-11
Creating a Local SPAN Session and Configuring Ingress Traffic	22-13
Specifying VLANs to Filter	22-15
Configuring RSPAN	22-16
RSPAN Configuration Guidelines	22-16
Configuring a VLAN as an RSPAN VLAN	22-17
Creating an RSPAN Source Session	22-18
Creating an RSPAN Destination Session	22-19
Creating an RSPAN Destination Session and Configuring Ingress Traffic	22-20
Specifying VLANs to Filter	22-22
Displaying SPAN and RSPAN Status	22-23

CHAPTER 23

Configuring RMON 23-1

Understanding RMON	23-1
Configuring RMON	23-2
Default RMON Configuration	23-3
Configuring RMON Alarms and Events	23-3
Collecting Group History Statistics on an Interface	23-5
Collecting Group Ethernet Statistics on an Interface	23-6
Displaying RMON Status	23-6

CHAPTER 24

Configuring System Message Logging 24-1

Understanding System Message Logging	24-1
Configuring System Message Logging	24-2
System Log Message Format	24-2
Default System Message Logging Configuration	24-3
Disabling Message Logging	24-4
Setting the Message Display Destination Device	24-4
Synchronizing Log Messages	24-5
Enabling and Disabling Timestamps on Log Messages	24-7
Enabling and Disabling Sequence Numbers in Log Messages	24-7
Defining the Message Severity Level	24-8
Limiting Syslog Messages Sent to the History Table and to SNMP	24-9
Configuring UNIX Syslog Servers	24-10
Logging Messages to a UNIX Syslog Daemon	24-10
Configuring the UNIX System Logging Facility	24-11
Displaying the Logging Configuration	24-12

CHAPTER 25

Configuring SNMP 25-1

- Understanding SNMP 25-1
 - SNMP Versions 25-2
 - SNMP Manager Functions 25-3
 - SNMP Agent Functions 25-4
 - SNMP Community Strings 25-4
 - Using SNMP to Access MIB Variables 25-5
 - SNMP Notifications 25-5
 - SNMP ifIndex MIB Object Values 25-6
- Configuring SNMP 25-6
 - Default SNMP Configuration 25-7
 - SNMP Configuration Guidelines 25-7
 - Disabling the SNMP Agent 25-8
 - Configuring Community Strings 25-8
 - Configuring SNMP Groups and Users 25-9
 - Configuring SNMP Notifications 25-11
 - Setting the Agent Contact and Location Information 25-14
 - Limiting TFTP Servers Used Through SNMP 25-15
 - SNMP Examples 25-15
- Displaying SNMP Status 25-16

CHAPTER 26

Configuring Network Security with ACLs 26-1

- Understanding ACLs 26-1
 - Supported ACLs 26-2
 - Port ACLs 26-2
 - VLAN Maps 26-3
 - Handling Fragmented and Unfragmented Traffic 26-4
- Configuring IP ACLs 26-5
 - Creating Standard and Extended IP ACLs 26-6
 - Access List Numbers 26-6
 - Creating a Numbered Standard ACL 26-7
 - Creating a Numbered Extended ACL 26-8
 - Creating Named Standard and Extended ACLs 26-11
 - Using Time Ranges with ACLs 26-13
 - Including Comments in ACLs 26-15
 - Applying an IP ACL to a Terminal Line 26-16
 - Applying an IP ACL to an Interface 26-16
 - Hardware and Software Treatment of IP ACLs 26-17

IP ACL Configuration Examples	26-17
Numbered ACLs	26-18
Extended ACLs	26-18
Named ACLs	26-18
Time Range Applied to an IP ACL	26-19
Commented IP ACL Entries	26-19
Creating Named MAC Extended ACLs	26-20
Applying a MAC ACL to a Layer 2 Interface	26-21
Configuring VLAN Maps	26-22
VLAN Map Configuration Guidelines	26-23
Creating a VLAN Map	26-23
Examples of ACLs and VLAN Maps	26-24
Applying a VLAN Map to a VLAN	26-26
Using VLAN Maps in Your Network	26-26
Wiring Closet Configuration	26-26
Denying Access to a Server on a VLAN	26-28
Displaying ACL Configuration	26-29
CHAPTER 27	Configuring QoS 27-1
Understanding QoS	27-1
Basic QoS Model	27-3
Classification	27-4
Classification Based on QoS ACLs	27-7
Classification Based on Class Maps and Policy Maps	27-7
Policing and Marking	27-8
Mapping Tables	27-10
Queueing and Scheduling Overview	27-11
Weighted Tail Drop	27-11
SRR Shaping and Sharing	27-12
Queueing and Scheduling on Ingress Queues	27-13
Queueing and Scheduling on Egress Queues	27-15
Packet Modification	27-17
Configuring Auto-QoS	27-18
Generated Auto-QoS Configuration	27-18
Effects of Auto-QoS on the Configuration	27-22
Auto-QoS Configuration Guidelines	27-22
Enabling Auto-QoS for VoIP	27-23
Auto-QoS Configuration Example	27-24
Displaying Auto-QoS Information	27-26

Configuring Standard QoS	27-26
Default Standard QoS Configuration	27-27
Default Ingress Queue Configuration	27-27
Default Egress Queue Configuration	27-28
Default Mapping Table Configuration	27-28
Standard QoS Configuration Guidelines	27-29
Enabling QoS Globally	27-30
Configuring Classification Using Port Trust States	27-30
Configuring the Trust State on Ports within the QoS Domain	27-31
Configuring the CoS Value for an Interface	27-33
Configuring a Trusted Boundary to Ensure Port Security	27-34
Configuring the DSCP Trust State on a Port Bordering Another QoS Domain	27-35
Configuring a QoS Policy	27-36
Classifying Traffic by Using ACLs	27-37
Classifying Traffic by Using Class Maps	27-40
Classifying, Policing, and Marking Traffic by Using Policy Maps	27-42
Classifying, Policing, and Marking Traffic by Using Aggregate Policers	27-45
Configuring DSCP Maps	27-47
Configuring the CoS-to-DSCP Map	27-47
Configuring the IP-Precedence-to-DSCP Map	27-48
Configuring the Policed-DSCP Map	27-49
Configuring the DSCP-to-CoS Map	27-50
Configuring the DSCP-to-DSCP-Mutation Map	27-51
Configuring Ingress Queue Characteristics	27-52
Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds	27-53
Allocating Buffer Space Between the Ingress Queues	27-54
Allocating Bandwidth Between the Ingress Queues	27-55
Configuring the Ingress Priority Queue	27-56
Configuring Egress Queue Characteristics	27-57
Configuration Guidelines	27-57
Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set	27-58
Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID	27-60
Configuring SRR Shaped Weights on Egress Queues	27-61
Configuring SRR Shared Weights on Egress Queues	27-62
Configuring the Egress Expedite Queue	27-63
Limiting the Bandwidth on an Egress Interface	27-64
Displaying Standard QoS Information	27-65

CHAPTER 28

Configuring EtherChannels 28-1

- Understanding EtherChannels 28-1
 - EtherChannel Overview 28-2
 - Port-Channel Interfaces 28-3
 - Port Aggregation Protocol 28-3
 - PAgP Modes 28-4
 - PAgP Interaction with Other Features 28-5
 - Link Aggregation Control Protocol 28-5
 - LACP Modes 28-6
 - LACP Interaction with Other Features 28-6
 - Load Balancing and Forwarding Methods 28-6
- Configuring EtherChannels 28-8
 - Default EtherChannel Configuration 28-9
 - EtherChannel Configuration Guidelines 28-9
 - Configuring Layer 2 EtherChannels 28-10
 - Configuring EtherChannel Load Balancing 28-12
 - Configuring the PAgP Learn Method and Priority 28-13
 - Configuring LACP Hot-Standby Ports 28-15
 - Configuring the LACP System Priority 28-15
 - Configuring the LACP Port Priority 28-16
 - Displaying EtherChannel, PAgP, and LACP Status 28-17

CHAPTER 29

Troubleshooting 29-1

- Recovering from Corrupted Software By Using the XMODEM Protocol 29-2
- Recovering from a Lost or Forgotten Password 29-4
 - Procedure with Password Recovery Enabled 29-5
 - Procedure with Password Recovery Disabled 29-6
- Recovering from a Command Switch Failure 29-8
 - Replacing a Failed Command Switch with a Cluster Member 29-8
 - Replacing a Failed Command Switch with Another Switch 29-10
- Recovering from Lost Cluster Member Connectivity 29-11
- Preventing Autonegotiation Mismatches 29-12
- SFP Module Security and Identification 29-12
- Using Ping 29-13
 - Understanding Ping 29-13
 - Executing Ping 29-13

Using Layer 2 Traceroute	29-14
Understanding Layer 2 Traceroute	29-14
Usage Guidelines	29-14
Displaying the Physical Path	29-15
Using IP Traceroute	29-16
Understanding IP Traceroute	29-16
Executing IP Traceroute	29-16
Using Debug Commands	29-17
Enabling Debugging on a Specific Feature	29-18
Enabling All-System Diagnostics	29-18
Redirecting Debug and Error Message Output	29-19
Using the show platform forward Command	29-19
Using the crashinfo File	29-21

APPENDIX A

Supported MIBs A-1

MIB List	A-1
Using FTP to Access the MIB Files	A-3

APPENDIX B

Working with the Cisco IOS File System, Configuration Files, and Software Images B-1

Working with the Flash File System	B-1
Displaying Available File Systems	B-2
Setting the Default File System	B-3
Displaying Information about Files on a File System	B-3
Changing Directories and Displaying the Working Directory	B-3
Creating and Removing Directories	B-4
Copying Files	B-4
Deleting Files	B-5
Creating, Displaying, and Extracting tar Files	B-5
Creating a tar File	B-6
Displaying the Contents of a tar File	B-6
Extracting a tar File	B-7
Displaying the Contents of a File	B-7
Working with Configuration Files	B-8
Guidelines for Creating and Using Configuration Files	B-9
Configuration File Types and Location	B-9
Creating a Configuration File By Using a Text Editor	B-10

Copying Configuration Files By Using TFTP	B-10
Preparing to Download or Upload a Configuration File By Using TFTP	B-10
Downloading the Configuration File By Using TFTP	B-11
Uploading the Configuration File By Using TFTP	B-11
Copying Configuration Files By Using FTP	B-12
Preparing to Download or Upload a Configuration File By Using FTP	B-13
Downloading a Configuration File By Using FTP	B-13
Uploading a Configuration File By Using FTP	B-15
Copying Configuration Files By Using RCP	B-16
Preparing to Download or Upload a Configuration File By Using RCP	B-16
Downloading a Configuration File By Using RCP	B-17
Uploading a Configuration File By Using RCP	B-18
Clearing Configuration Information	B-19
Clearing the Startup Configuration File	B-19
Deleting a Stored Configuration File	B-19
Working with Software Images	B-20
Image Location on the Switch	B-20
tar File Format of Images on a Server or Cisco.com	B-21
Copying Image Files By Using TFTP	B-22
Preparing to Download or Upload an Image File By Using TFTP	B-22
Downloading an Image File By Using TFTP	B-23
Uploading an Image File By Using TFTP	B-24
Copying Image Files By Using FTP	B-25
Preparing to Download or Upload an Image File By Using FTP	B-25
Downloading an Image File By Using FTP	B-26
Uploading an Image File By Using FTP	B-28
Copying Image Files By Using RCP	B-29
Preparing to Download or Upload an Image File By Using RCP	B-29
Downloading an Image File By Using RCP	B-31
Uploading an Image File By Using RCP	B-33

APPENDIX C

Unsupported Commands in Cisco IOS Release 12.1(19)EA1 C-1

Access Control Lists	C-1
Unsupported Privileged EXEC Commands	C-1
Unsupported Global Configuration Commands	C-1
Unsupported Debug Commands	C-1
IGMP Snooping Commands	C-2
Unsupported Global Configuration Commands	C-2

Interface Commands	C-2
Unsupported Privileged EXEC Commands	C-2
Unsupported Global Configuration Commands	C-2
Unsupported Interface Configuration Commands	C-2
Network Address Translation (NAT) Commands	C-2
Unsupported User EXEC Commands	C-2
Unsupported Global Configuration Commands	C-3
Unsupported Interface Configuration Commands	C-3
RADIUS	C-3
Unsupported Global Configuration Commands	C-3
SNMP	C-3
Unsupported Global Configuration Commands	C-3
Spanning Tree	C-3
Unsupported Global Configuration Commands	C-3
Unsupported Interface Configuration Commands	C-4
VLAN	C-4
Unsupported vlan-config Commands	C-4
Unsupported User EXEC Commands	C-4
VTP	C-4
Unsupported Privileged EXEC Commands	C-4
Miscellaneous	C-4
Unsupported Global Configuration Commands	C-4



Preface

Audience

This guide is for the networking professional managing the Catalyst 2970 switch, hereafter referred to as the *switch*. Before using this guide, you should have experience working with the Cisco IOS software and be familiar with the concepts and terminology of Ethernet and local area networking.

Purpose

This guide provides the information that you need to configure features on your switch. The Catalyst 2970 enhanced software image provides enterprise-class intelligent services such as access control lists (ACLs) and quality of service (QoS) features.

This guide provides procedures for using the commands that have been created or changed for use with the Catalyst 2970 switch. It does not provide detailed information about these commands. For detailed information about these commands, refer to the *Catalyst 2970 Switch Command Reference* for this release. For information about the standard Cisco IOS Release 12.1 commands, refer to the Cisco IOS documentation set available from the Cisco.com home page at **Service and Support > Technical Documents**. On the Cisco Product Documentation home page, select **Release 12.1** from the Cisco IOS Software drop-down list.

This guide also includes an overview of the Cluster Management Suite (CMS), a web-based switch management interface that helps you create and manage clusters of switches. This guide does not provide field-level descriptions of the CMS windows nor does it provide the procedures for configuring switches and switch clusters from CMS. For all CMS window descriptions and procedures, refer to the CMS online help, which is integrated with the software image.

This guide does not describe system messages you might encounter or how to install your switch. For more information, refer to the *Catalyst 2970 Switch System Message Guide* for this release and to the *Catalyst 2970 Switch Hardware Installation Guide*.

Conventions

This publication uses these conventions to convey instructions and information:

Command descriptions use these conventions:

- Commands and keywords are in **boldface** text.
- Arguments for which you supply values are in *italic*.
- Square brackets ([]) mean optional elements.
- Braces ({ }) group required choices, and vertical bars (|) separate the alternative elements.
- Braces and vertical bars within square brackets ([{ | }]) mean a required choice within an optional element.

Interactive examples use these conventions:

- Terminal sessions and system displays are in `screen` font.
- Information you enter is in **boldface screen** font.
- Nonprinting characters, such as passwords or tabs, are in angle brackets (<>).

Notes, cautions, and timesavers use these conventions and symbols:



Note

Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this manual.



Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.



Timesaver

Means the following *will help you solve a problem*. The tips information might not be troubleshooting or even an action, but could be useful information.

Related Publications

These documents provide complete information about the switch and are available from this Cisco.com site:

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat2970/index.htm>

You can order printed copies of documents with a DOC-xxxxxx= number from the Cisco.com sites and from the telephone numbers listed in the “[Obtaining Documentation](#)” section on page xxix.

- *Release Notes for the Catalyst 2970 Switch* (not orderable but available on Cisco.com)
- *Catalyst 2970 Switch Software Configuration Guide* (order number DOC-7815462=)
- *Catalyst 2970 Switch Command Reference* (order number DOC-7815464=)
- *Catalyst 2970 Switch System Message Guide* (order number DOC-7815465=)
- Cluster Management Suite (CMS) online help (available only from the switch CMS software)
- *Catalyst 2970 Switch Hardware Installation Guide* (order number DOC-7815469=)
- *Cisco CWDM GBIC and CWDM SFP Installation Note* (not orderable but available on Cisco.com)

Obtaining Documentation

Cisco provides several ways to obtain documentation, technical assistance, and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

Cisco.com

You can access the most current Cisco documentation on the World Wide Web at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

International Cisco websites can be accessed from this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which may have shipped with your product. The Documentation CD-ROM is updated regularly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual or quarterly subscription.

Registered Cisco.com users can order a single Documentation CD-ROM (product number DOC-CONDOCCD=) through the Cisco Ordering tool:

http://www.cisco.com/en/US/partner/ordering/ordering_place_order_ordering_tool_launch.html

All users can order annual or quarterly subscriptions through the online Subscription Store:

<http://www.cisco.com/go/subscription>

Ordering Documentation

You can find instructions for ordering documentation at this URL:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

<http://www.cisco.com/en/US/partner/ordering/index.shtml>

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. On the Cisco Documentation home page, click **Feedback** at the top of the page.

You can send your comments in e-mail to bug-doc@cisco.com.

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Customer Document Ordering
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, the Cisco Technical Assistance Center (TAC) provides 24-hour, award-winning technical support services, online and over the phone. Cisco.com features the Cisco TAC website as an online starting point for technical assistance.

Cisco TAC Website

The Cisco TAC website (<http://www.cisco.com/tac>) provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The Cisco TAC website is available 24 hours a day, 365 days a year.

Accessing all the tools on the Cisco TAC website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a login ID or password, register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

Opening a TAC Case

The online TAC Case Open Tool (<http://www.cisco.com/tac/caseopen>) is the fastest way to open P3 and P4 cases. (Your network is minimally impaired or you require product information). After you describe your situation, the TAC Case Open Tool automatically recommends resources for an immediate solution. If your issue is not resolved using these recommendations, your case will be assigned to a Cisco TAC engineer.

For P1 or P2 cases (your production network is down or severely degraded) or if you do not have Internet access, contact Cisco TAC by telephone. Cisco TAC engineers are assigned immediately to P1 and P2 cases to help keep your business operations running smoothly.

To open a case by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete listing of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

TAC Case Priority Definitions

To ensure that all cases are reported in a standard format, Cisco has established case priority definitions.

Priority 1 (P1)—Your network is “down” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

Priority 2 (P2)—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

Priority 3 (P3)—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

Priority 4 (P4)—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- The Cisco Product Catalog describes the networking products offered by Cisco Systems, as well as ordering and customer support services. Access the Cisco Product Catalog at this URL:

http://www.cisco.com/en/US/products/products_catalog_links_launch.html

- Cisco Press publishes a wide range of networking publications. Cisco suggests these titles for new and experienced users: Internetworking Terms and Acronyms Dictionary, Internetworking Technology Handbook, Internetworking Troubleshooting Guide, and the Internetworking Design Guide. For current Cisco Press titles and other information, go to Cisco Press online at this URL:

<http://www.ciscopress.com>

- Packet magazine is the Cisco quarterly publication that provides the latest networking trends, technology breakthroughs, and Cisco products and solutions to help industry professionals get the most from their networking investment. Included are networking deployment and troubleshooting tips, configuration examples, customer case studies, tutorials and training, certification information, and links to numerous in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/go/packet>

- iQ Magazine is the Cisco bimonthly publication that delivers the latest information about Internet business strategies for executives. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- Internet Protocol Journal is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

http://www.cisco.com/en/US/about/ac123/ac147/about_cisco_the_internet_protocol_journal.html

- Training—Cisco offers world-class networking training. Current offerings in network training are listed at this URL:

<http://www.cisco.com/en/US/learning/index.html>



Overview

This chapter provides these topics about the Catalyst 2970 switch software:

- [Features, page 1-1](#)
- [Default Settings After Initial Switch Configuration, page 1-8](#)
- [Network Configuration Examples, page 1-10](#)
- [Where to Go Next, page 1-15](#)

In this document, IP refers to IP version 4 (IPv4).

Features

Some features noted in this chapter are available only on the cryptographic (that is, supports encryption) version of the switch software image. You must obtain authorization to use this feature and to download the cryptographic version of the software from Cisco.com. For more information, refer to the release notes for this release.

The Catalyst 2970 switches have these features:

- [Ease-of-Use and Ease-of-Deployment Features, page 1-2](#)
- [Performance Features, page 1-3](#)
- [Management Options, page 1-3](#)
- [Manageability Features, page 1-4](#) (includes a feature requiring the cryptographic [that is, supports encryption] version of the switch software image)
- [Availability Features, page 1-4](#)
- [VLAN Features, page 1-5](#)
- [Security Features, page 1-5](#) (includes a feature requiring the cryptographic [that is, supports encryption] version of the switch software image)
- [QoS and CoS Features, page 1-6](#)
- [Monitoring Features, page 1-7](#)

Ease-of-Use and Ease-of-Deployment Features

- Express Setup for quickly configuring a switch for the first time with basic IP information, contact information, switch and Telnet passwords, and Simple Network Management Protocol (SNMP) information through a browser-based program
- User-defined SmartPort macros for creating custom switch configurations for simplified deployment across the network
- Cluster Management Suite (CMS) graphical user interface (GUI) for
 - Simplifying and minimizing switch and switch cluster management through a supported web browser from anywhere in your intranet.
 - Accomplishing multiple configuration tasks from a single CMS window without needing to remember command-line interface (CLI) commands to accomplish specific tasks.
 - Interactive guide mode that guides you in configuring complex features such as VLANs, ACLs, and quality of service (QoS).
 - Automated configuration wizards that prompt you to provide only the minimum required information to configure complex features such as QoS priorities for video traffic, priority levels for data applications, and security.
 - Applying actions to multiple ports and multiple switches at the same time, such as VLAN and QoS settings, inventory and statistic reports, link- and switch-level monitoring and troubleshooting, and multiple switch software upgrades.
 - Viewing a topology of interconnected devices to identify existing switch clusters and eligible switches that can join a cluster and to identify link information between switches.
 - Monitoring real-time status of a switch or multiple switches from the LEDs on the front-panel images. The system, redundant power system (RPS), and port LED colors on the images are similar to those used on the physical LEDs.
- Switch clustering technology for
 - Unified configuration, monitoring, authentication, and software upgrade of multiple, cluster-capable switches, regardless of their geographic proximity and interconnection media, including Ethernet, Fast Ethernet, Fast EtherChannel, small form-factor pluggable (SFP) modules, Gigabit Ethernet, and Gigabit EtherChannel connections. Refer to the release notes for a list of cluster-capable switches.

- Automatic discovery of candidate switches and creation of clusters of up to 16 switches that can be managed through a single IP address.
- Extended discovery of cluster candidates that are not directly connected to the command switch.

Performance Features

- Autosensing of port speed and autonegotiation of duplex mode on all switch ports for optimizing bandwidth
- Automatic-medium-dependent interface crossover (Auto-MDIX) capability on 10/100/1000 Mbps interfaces that enables the interface to automatically detect the required cable connection type (straight through or crossover) and configure the connection appropriately
- IEEE 802.3X flow control on all ports (the switch does not send pause frames)
- EtherChannel for enhanced fault tolerance and for providing up to 8 Gbps (Gigabit EtherChannel) or 800 Mbps (Fast EtherChannel) full duplex of bandwidth between switches, routers, and servers
- Port Aggregation Protocol (PAgP) and Link Aggregation Control Protocol (LACP) for automatic creation of EtherChannel links
- Forwarding of Layer 2 packets at Gigabit line rate
- Per-port storm control for preventing broadcast, multicast, and unicast storms
- Port blocking on forwarding unknown Layer 2 unknown unicast, multicast, and bridged broadcast traffic
- Internet Group Management Protocol (IGMP) snooping for IGMP versions 1, 2, and 3 for efficiently forwarding multimedia and multicast traffic
- IGMP report suppression for sending only one IGMP report per multicast router query to the multicast devices (supported only for IGMPv1 or IGMPv2 queries)
- Multicast VLAN registration (MVR) to continuously send multicast streams in a multicast VLAN while isolating the streams from subscriber VLANs for bandwidth and security reasons
- IGMP filtering for controlling the set of multicast groups to which hosts on a switch port can belong
- IGMP throttling for configuring the action when the maximum number of entries is in the IGMP forwarding table

Management Options

- CMS—CMS is a graphical user interface that can be launched from anywhere in your network through a web browser such as Netscape Communicator or Microsoft Internet Explorer. CMS is already installed on the switch. For more information about CMS, see [Chapter 3, “Getting Started with CMS.”](#)
- CLI—The Cisco IOS CLI software is enhanced to support desktop- and multilayer-switching features. You can access the CLI either by connecting your management station directly to the switch console port or by using Telnet from a remote management station. For more information about the CLI, see [Chapter 2, “Using the Command-Line Interface.”](#)
- SNMP—SNMP management applications such as CiscoWorks2000 LAN Management Suite (LMS) and HP OpenView. You can manage from an SNMP-compatible management station that is running platforms such as HP OpenView or SunNet Manager. The switch supports a comprehensive set of MIB extensions and four remote monitoring (RMON) groups. For more information about using SNMP, see [Chapter 25, “Configuring SNMP.”](#)

Manageability Features

**Note**

The encrypted Secure Shell (SSH) feature listed in this section is available only on the cryptographic (that is, supports encryption) version of the switch software image.

- Dynamic Host Configuration Protocol (DHCP) for automating configuration of switch information (such as IP address, default gateway, host name, and Domain Name System [DNS] and Trivial File Transfer Protocol (TFTP) server names)
- DHCP relay for forwarding User Datagram Protocol (UDP) broadcasts, including IP address requests, from DHCP clients
- DHCP server for automatic assignment of IP addresses and other DHCP options to IP hosts
- Directed unicast requests to a DNS server for identifying a switch through its IP address and its corresponding host name and to a TFTP server for administering software upgrades from a TFTP server
- Address Resolution Protocol (ARP) for identifying a switch through its IP address and its corresponding Media Access Control (MAC) address
- Unicast MAC address filtering to drop packets with specific source or destination MAC addresses
- Cisco Discovery Protocol (CDP) versions 1 and 2 for network topology discovery and mapping between the switch and other Cisco devices on the network
- Network Time Protocol (NTP) for providing a consistent timestamp to all switches from an external source
- Cisco IOS File System (IFS) for providing a single interface to all file systems that the switch uses
- In-band management access through CMS over a Netscape Communicator or Microsoft Internet Explorer browser session
- In-band management access for up to 16 simultaneous Telnet connections for multiple CLI-based sessions over the network
- In-band management access for up to five simultaneous, encrypted Secure Shell (SSH) connections for multiple CLI-based sessions over the network (requires the cryptographic [that is, supports encryption] version of the switch software image)
- In-band management access through SNMP versions 1 and 2c, and 3 get and set requests
- Out-of-band management access through the switch console port to a directly attached terminal or to a remote terminal through a serial connection or a modem

**Note**

For additional descriptions of the management interfaces, see the [“Network Configuration Examples” section on page 1-10](#).

Availability Features

- UniDirectional Link Detection (UDLD) and aggressive UDLD for detecting and disabling unidirectional links on fiber-optic interfaces caused by incorrect fiber-optic wiring or port faults
- IEEE 802.1D Spanning Tree Protocol (STP) for redundant backbone connections and loop-free networks. STP has these features:
 - Up to 128 spanning-tree instances supported
 - Per-VLAN spanning-tree plus (PVST+) for balancing load across VLANs

- Rapid PVST+ for balancing load across VLANs and providing rapid convergence of spanning-tree instances
- UplinkFast and BackboneFast for fast convergence after a spanning-tree topology change and for achieving load balancing between redundant uplinks, including Gigabit uplinks
- IEEE 802.1S Multiple Spanning Tree Protocol (MSTP) for grouping VLANs into a spanning-tree instance and for providing multiple forwarding paths for data traffic and load balancing and IEEE 802.1W Rapid Spanning Tree Protocol (RSTP) for rapid convergence of the spanning tree by immediately transitioning root and designated ports to the forwarding state
- Optional spanning-tree features available in PVST+, rapid-PVST+, and MSTP mode:
 - Port Fast for eliminating the forwarding delay by enabling a port to immediately transition from the blocking state to the forwarding state
 - BPDU guard for shutting down Port Fast-enabled ports that receive bridge protocol data units (BPDUs)
 - BPDU filtering for preventing a Port Fast-enabled port from sending or receiving BPDUs
 - Root guard for preventing switches outside the network core from becoming the spanning-tree root
 - Loop guard for preventing alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link
- RPS support through the Cisco RPS 300 and Cisco RPS 675 for enhancing power reliability

VLAN Features

- Support for up to 1005 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth
- Support for VLAN IDs in the full 1 to 4094 range allowed by the IEEE 802.1Q standard
- VLAN Query Protocol (VQP) for dynamic VLAN membership
- Inter-Switch Link (ISL) and IEEE 802.1Q trunking encapsulation on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources
- Dynamic Trunking Protocol (DTP) for negotiating trunking on a link between two devices and for negotiating the type of trunking encapsulation (802.1Q or ISL) to be used
- VLAN Trunking Protocol (VTP) and VTP pruning for reducing network traffic by restricting flooded traffic to links destined for stations receiving the traffic
- Voice VLAN for creating subnets for voice traffic from Cisco IP Phones
- VLAN1 minimization for reducing the risk of spanning-tree loops or storms by allowing VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received on the trunk. The switch CPU continues to send and receive control protocol frames.

Security Features



Note

The Kerberos feature listed in this section is available only on the cryptographic (that is, supports encryption) version of the switch software image.

- Password-protected access (read-only and read-write access) to management interfaces (CMS and CLI) for protection against unauthorized configuration changes

- Multilevel security for a choice of security level, notification, and resulting actions
- Static MAC addressing for ensuring security
- Protected port option for restricting the forwarding of traffic to designated ports on the same switch
- Port security option for limiting and identifying MAC addresses of the stations allowed to access the port
- Port security aging to set the aging time for secure addresses on a port
- BPDU guard for shutting down a Port Fast-configured port when an invalid configuration occurs
- Standard and extended IP access control lists (ACLs) for defining security policies in both directions on VLANs and inbound on Layer 2 interfaces (port ACLs)
- Extended MAC access control lists for defining security policies in the inbound direction on Layer 2 interfaces
- VLAN ACLs (VLAN maps) for providing intra-VLAN security by filtering traffic based on information in the MAC, IP, and TCP/User Datagram Protocol (UDP) headers
- Source and destination MAC-based ACLs for filtering non-IP traffic
- DHCP snooping to filter untrusted DHCP messages between untrusted hosts and DHCP servers
- IEEE 802.1X port-based authentication to prevent unauthorized devices (clients) from gaining access to the network
 - 802.1X with VLAN assignment for restricting 802.1X-authenticated users to a specified VLAN
 - 802.1X with port security for controlling access to 802.1X ports
 - 802.1X with voice VLAN to permit an IP phone access to the voice VLAN regardless of the authorized or unauthorized state of the port
 - 802.1X with guest VLAN to provide limited services to non-802.1X-compliant users
- Terminal Access Controller Access Control System Plus (TACACS+), a proprietary feature for managing network security through a TACACS server
- Remote Authentication Dial-In User Service (RADIUS) for verifying the identity of, granting access to, and tracking the actions of remote users through authentication, authorization, and accounting (AAA) services
- Kerberos security system to authenticate requests for network resources by using a trusted third party (requires the cryptographic [that is, supports encryption] version of the switch software image)

QoS and CoS Features

- Automatic QoS (auto-QoS) to simplify the deployment of existing QoS features by classifying traffic and configuring egress queues (voice over IP only)
- Classification
 - IP type-of-service/Differentiated Services Code Point (IP TOS/DSCP) and 802.1P CoS marking priorities on a per-port basis for protecting the performance of mission-critical applications
 - IP TOS/DSCP and 802.1P CoS marking based on flow-based packet classification (classification based on information in the MAC, IP, and TCP/UDP headers) for high-performance quality of service at the network edge, allowing for differentiated service levels for different types of network traffic and for prioritizing mission-critical traffic in the network
 - Trusted port states (CoS, DSCP, and IP precedence) within a QoS domain and with a port bordering another QoS domain

- Trusted boundary for detecting the presence of a Cisco IP phone, trusting the CoS value received, and ensuring port security
- Policing
 - Traffic-policing policies on the switch port for managing how much of the port bandwidth should be allocated to a specific traffic flow
 - Aggregate policing for policing traffic flows in aggregate to restrict specific applications or traffic flows to metered, predefined rates
- Out-of-Profile
 - Out-of-profile markdown for packets that exceed bandwidth utilization limits
- Ingress queueing and scheduling
 - Two configurable ingress queues for user traffic (one queue can be the priority queue)
 - Weighted tail drop (WTD) as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications
 - Shaped round robin (SRR) as the scheduling service for determining the rate at which packets are dequeued to the internal ring (sharing is the only supported mode on ingress queues)
- Egress queues and scheduling
 - Four egress queues per port
 - WTD as the congestion-avoidance mechanism for managing the queue lengths and providing drop precedences for different traffic classifications
 - SRR as the scheduling service for determining the rate at which packets are dequeued to the egress interface (shaping or sharing is supported on egress queues). Shaped egress queues are guaranteed but limited to using a share of port bandwidth. Shared egress queues are also guaranteed a configured share of bandwidth, but can use more than the guarantee if other queues become empty and do not use their share of the bandwidth.

Monitoring Features

- Switch LEDs that provide port- and switch-level status
- MAC address notification traps and RADIUS accounting for tracking users on a network by storing the MAC addresses that the switch has learned or removed
- Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) for traffic monitoring on any port or VLAN
- SPAN and RSPAN support of Intrusion Detection Systems (IDS) to monitor, repel, and report network security violations
- Four groups (history, statistics, alarms, and events) of embedded RMON agents for network monitoring and traffic analysis
- Syslog facility for logging system messages about authentication or authorization errors, resource issues, and time-out events
- Layer 2 traceroute to identify the physical path that a packet takes from a source device to a destination device
- Time Domain Reflector (TDR) to diagnose and resolve cabling problems on copper Ethernet 10/100/1000 ports

Default Settings After Initial Switch Configuration

The switch is designed for plug-and-play operation, requiring only that you assign basic IP information to the switch and connect it to the other devices in your network. If you have specific network needs, you can change the interface-specific and system-wide settings.

If you do not configure the switch at all, the switch operates with the default settings listed in [Table 1-1](#). This table lists the key software features, their defaults, and where to find more information about the features.

For information about setting up the initial switch configuration and assigning basic IP information to the switch, refer to the hardware installation guide.

Table 1-1 Default Settings After Initial Switch Configuration

Feature	Default Setting	More information in...
Switch IP address, subnet mask, and default gateway	0.0.0.0	Chapter 4, “Assigning the Switch IP Address and Default Gateway”
Domain name	None	
DHCP	DHCP client enabled	
Switch cluster	Disabled	Chapter 5, “Clustering Switches”
Passwords	None defined	Chapter 6, “Administering the Switch”
TACACS+	Disabled	
RADIUS	Disabled	
System name and prompt	<i>Switch</i>	
NTP	Enabled	
DNS	Enabled	Chapter 8, “Configuring 802.1X Port-Based Authentication”
802.1X	Disabled	
Port parameters		
Operating mode	Layer 2 (switchport)	Chapter 9, “Configuring Interface Characteristics”
Interface speed and duplex mode	Autonegotiate	
Auto-MDIX	Disabled	
Flow control	Off	
SmartPort macros	None defined	Chapter 10, “Configuring SmartPort Macros”
VLANs		
Default VLAN	VLAN 1	Chapter 11, “Configuring VLANs”
VLAN trunking	Dynamic auto (DTP)	
Trunk encapsulation	Negotiate	
VTP mode	Server	Chapter 12, “Configuring VTP”
VTP version	1	
Voice VLAN	Disabled	Chapter 13, “Configuring Voice VLAN”
STP	PVST+ enabled on VLAN 1	Chapter 14, “Configuring STP”

Table 1-1 *Default Settings After Initial Switch Configuration (continued)*

Feature	Default Setting	More information in...
MSTP	Disabled	Chapter 15, “Configuring MSTP”
Optional spanning-tree features	Disabled	Chapter 16, “Configuring Optional Spanning-Tree Features”
DHCP snooping		
DHCP snooping	Disabled	Chapter 17, “Configuring DHCP Features”
DHCP snooping information option	Enabled	
IGMP snooping		
IGMP snooping	Enabled	Chapter 18, “Configuring IGMP Snooping and MVR”
IGMP filters	None applied	
IGMP throttling	Deny	
MVR	Disabled	
Port-based Traffic		
Broadcast, multicast, and unicast storm control	Disabled	Chapter 19, “Configuring Port-Based Traffic Control”
Protected ports	None defined	
Unicast and multicast traffic flooding	Not blocked	
Secure ports	None configured	
CDP	Enabled	Chapter 20, “Configuring CDP”
UDLD	Disabled	Chapter 21, “Configuring UDLD”
SPAN and RSPAN	Disabled	Chapter 22, “Configuring SPAN and RSPAN”
RMON	Disabled	Chapter 23, “Configuring RMON”
Syslog messages	Enabled; displayed on the console	Chapter 24, “Configuring System Message Logging”
SNMP	Enabled; version 1	Chapter 25, “Configuring SNMP”
ACLs	None configured	Chapter 26, “Configuring Network Security with ACLs”
QoS	Disabled	Chapter 27, “Configuring QoS”
EtherChannels	None configured	Chapter 28, “Configuring EtherChannels”

Network Configuration Examples

This section provides network configuration concepts and includes examples of using the switch to create dedicated network segments and interconnecting the segments through Gigabit Ethernet connections.

- [“Design Concepts for Using the Switch” section on page 1-10](#)
- [“Small to Medium-Sized Network Using Catalyst 2970 Switches” section on page 1-14](#)
- [“Long-Distance, High-Bandwidth Transport Configuration” section on page 1-15](#)

Design Concepts for Using the Switch

As your network users compete for network bandwidth, it takes longer to send and receive data. When you configure your network, consider the bandwidth required by your network users and the relative priority of the network applications they use.

[Table 1-2](#) describes what can cause network performance to degrade and how you can configure your network to increase the bandwidth available to your network users.

Table 1-2 *Increasing Network Performance*

Network Demands	Suggested Design Methods
Too many users on a single network segment and a growing number of users accessing the Internet	<ul style="list-style-type: none"> • Create smaller network segments so that fewer users share the bandwidth, and use VLANs and IP subnets to place the network resources in the same logical network as the users who access those resources most. • Use full-duplex operation between the switch and its connected workstations.
<ul style="list-style-type: none"> • Increased power of new PCs, workstations, and servers • High bandwidth demand from networked applications (such as e-mail with large attached files) and from bandwidth-intensive applications (such as multimedia) 	<ul style="list-style-type: none"> • Connect global resources—such as servers and routers to which the network users require equal access—directly to the high-speed switch ports so that they have their own high-speed segment. • Use the EtherChannel feature between the switch and its connected servers and routers.

Bandwidth alone is not the only consideration when designing your network. As your network traffic profiles evolve, consider providing network services that can support applications for voice and data integration, multimedia integration, application prioritization, and security. [Table 1-3](#) describes some network demands and how you can meet those demands.

Table 1-3 Providing Network Services

Network Demands	Suggested Design Methods
Efficient bandwidth usage for multimedia applications and guaranteed bandwidth for critical applications	<ul style="list-style-type: none"> • Use IGMP snooping to efficiently forward multimedia and multicast traffic. • Use other QoS mechanisms such as packet classification, marking, scheduling, and congestion avoidance to classify traffic with the appropriate priority level, thereby providing maximum flexibility and support for mission-critical, unicast, and multicast and multimedia applications. • Use MVR to continuously send multicast streams in a multicast VLAN but to isolate the streams from subscriber VLANs for bandwidth and security reasons.
High demand on network redundancy and availability to provide <i>always on</i> mission-critical applications	<ul style="list-style-type: none"> • Use Hot Standby Router Protocol (HSRP) for cluster command switch redundancy. • Use VLAN trunks and BackboneFast for traffic-load balancing on the uplink ports so that the uplink port with a lower relative port cost is selected to carry the VLAN traffic.
An evolving demand for IP telephony	<ul style="list-style-type: none"> • Use QoS to prioritize applications such as IP telephony during congestion and to help control both delay and jitter within the network. • Use switches that support at least two queues per port to prioritize voice and data traffic as either high- or low-priority, based on 802.1P/Q. The Catalyst 2970 switch supports at least four queues per port. • Use voice VLAN IDs (VVIDs) to provide separate VLANs for voice traffic.
A growing demand for using existing infrastructure to transport data and voice from a home or office to the Internet or an intranet at higher speeds	<p>Use the Catalyst Long-Reach Ethernet (LRE) switches to provide up to 15 Mb of IP connectivity over existing infrastructure, such as existing telephone lines.</p> <p>Note LRE is the technology used in the Catalyst 2900 LRE XL and Catalyst 2950 LRE switches. Refer to the documentation sets specific to these switches for LRE information.</p>

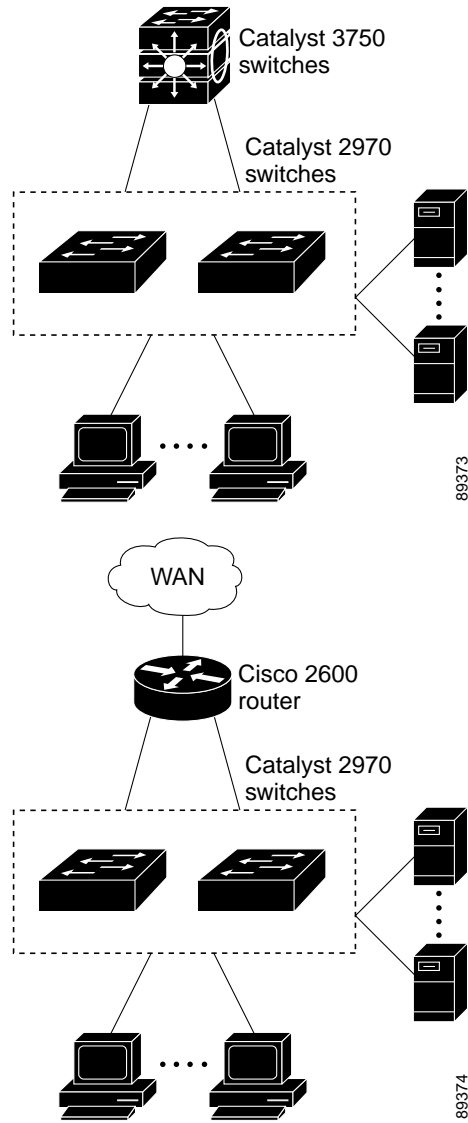
You can use the switches to create the following:

- Cost-effective Gigabit-to-the-desktop for high-performance workgroups ([Figure 1-1](#))—For high-speed access to network resources, you can use Catalyst 2970 switches in the access layer to provide Gigabit Ethernet to the desktop. To prevent congestion, use QoS DSCP marking priorities on these switches. For high-speed IP forwarding at the distribution layer, connect the switches in the access layer to a Gigabit multilayer switch with routing capability, such as a Catalyst 3750 switch, or to a router

The first illustration is of an isolated high-performance workgroup, where the Catalyst 2970 switches are connected to Catalyst 3750 switches in the distribution layer. The second illustration is of a high-performance workgroup in a branch office, where the Catalyst 2970 switches are connected to a router in the distribution layer.

Each switch in this configuration provides users with a dedicated 1-Gbps connection to network resources. Using SFP modules also provides flexibility in media and distance options through fiber-optic connections.

Figure 1-1 High-Performance Workgroup (Gigabit-to-the-Desktop)



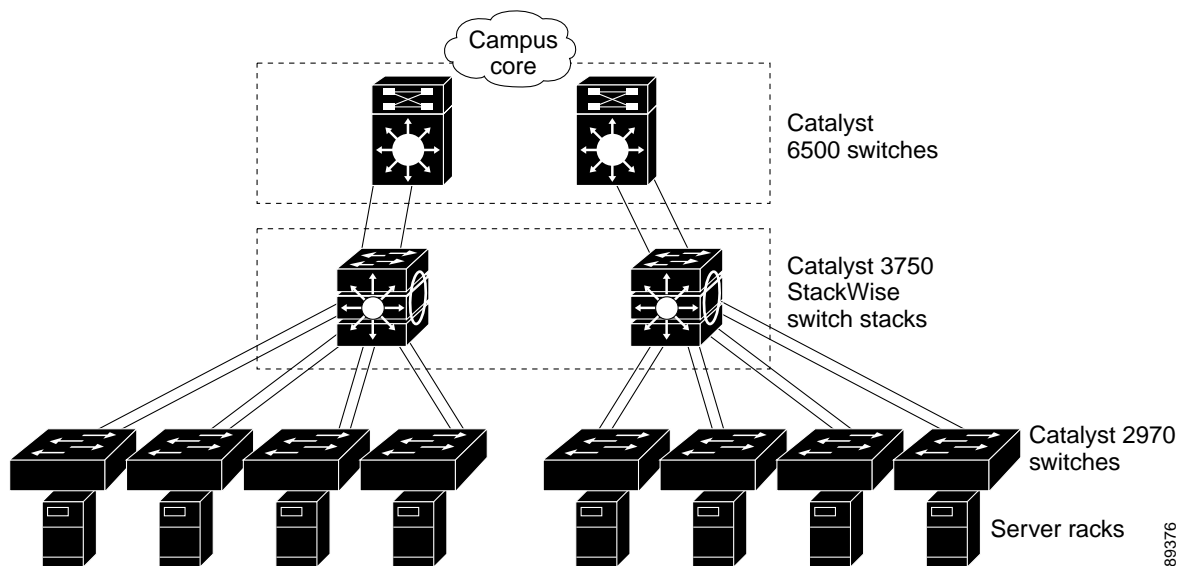
- Server aggregation ([Figure 1-2](#))—You can use the switches to interconnect groups of servers, centralizing physical security and administration of your network. For high-speed IP forwarding at the distribution layer, connect the switches in the access layer to multilayer switches with routing capability. The Gigabit interconnections minimize latency in the data flow.

QoS and policing on the switches provide preferential treatment for certain data streams, if required. They segment traffic streams into different paths for processing. Security features on the switch ensure rapid handling of packets.

Dual homing of servers to the switches with redundant Gigabit EtherChannel provides fault tolerance from the server racks to the core.

Using dual SFP uplinks from the Catalyst 2970 switches provide redundant uplinks to the network core. Using SFP modules provides flexibility in media and distance options through fiber-optic connections.

Figure 1-2 Server Aggregation



89376

Small to Medium-Sized Network Using Catalyst 2970 Switches

Figure 1-3 shows a configuration for a network of up to 500 employees. This network uses Catalyst 2970 switches with high-speed connections to two routers. For network reliability and load balancing, this network has HSRP enabled on the routers. This ensures connectivity to the Internet, WAN, and mission-critical network resources in case one of the routers fails. The switches are using EtherChannel for load sharing.

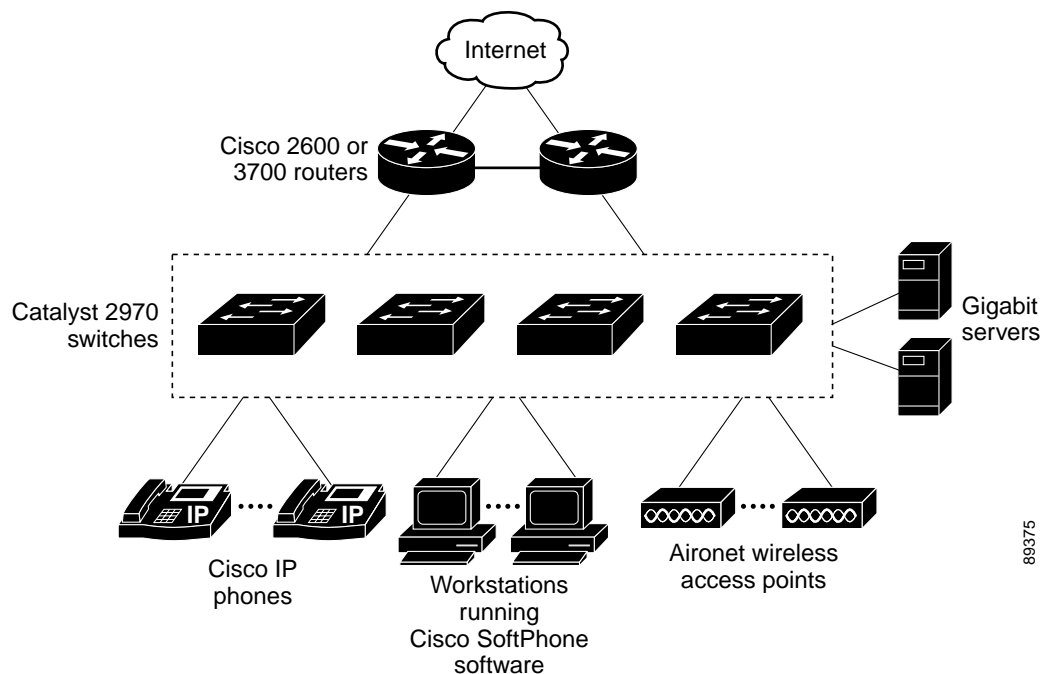
The switches are connected to workstations, Cisco IP Phones, and local servers. The switches are interconnected through Gigabit interfaces. This network uses VLANs to logically segment the network into well-defined broadcast groups and for security management. Data and multimedia traffic are configured on the same VLAN. Voice traffic from the Cisco IP Phones are configured on separate VVIDs. If data, multimedia, and voice traffic are assigned to the same VLAN, only one VLAN can be configured per wiring closet. For any switch port connected to Cisco IP Phones, 802.1P/Q QoS gives voice traffic forwarding-priority over data traffic. Cisco IP Phones not connected to Catalyst Power over Ethernet (PoE) switches must be connected to AC power sources to receive power.

When an end station in one VLAN needs to communicate with an end station in another VLAN, a router routes the traffic to the appropriate destination VLAN. In this network, the routers are providing inter-VLAN routing. VLAN access control lists (VLAN maps) provide intra-VLAN security and prevent unauthorized users from accessing critical pieces of the network.

In addition to inter-VLAN routing, the routers provide QoS mechanisms such as DSCP priorities to prioritize the different types of network traffic and to deliver high-priority traffic in a predictable manner. If congestion occurs, QoS drops low-priority traffic to allow delivery of high-priority traffic.

The routers also provide firewall services, Network Address Translation (NAT) services, voice-over-IP (VoIP) gateway services, and WAN and Internet access.

Figure 1-3 Catalyst 2970 Switches in a Collapsed Backbone Configuration



89375

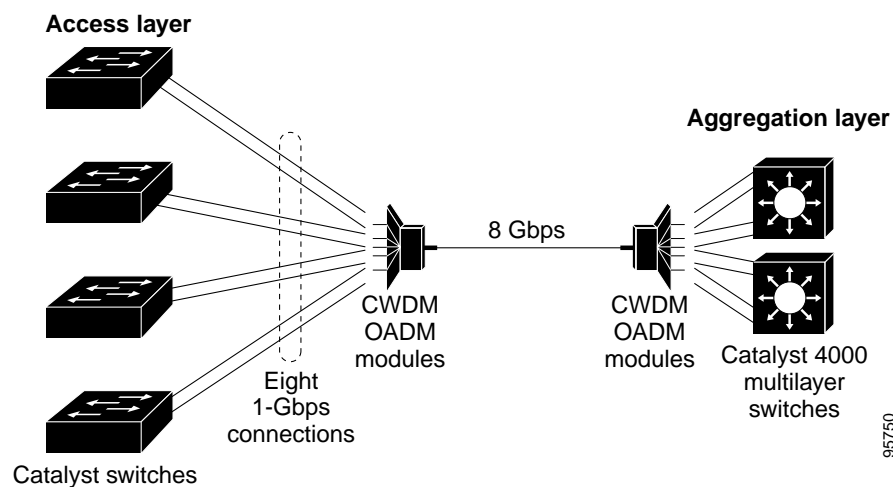
Long-Distance, High-Bandwidth Transport Configuration

Figure 1-4 shows a configuration for transporting 8 Gigabits of data over a single fiber-optic cable. The Catalyst switches have Coarse Wave Division Multiplexer (CWDM) fiber-optic SFP modules installed. Depending on the CWDM SFP module, data is sent at wavelengths from 1470 nm to 1610 nm. The higher the wavelength, the farther the transmission can travel. A common wavelength used for long-distance transmissions is 1550 nm.

The CWDM SFP modules connect to CWDM optical add/drop multiplexer (OADM) modules over distances of up to 393,701 feet (74.5 miles or 120 km). The CWDM OADM modules combine (or *multiplex*) the different CWDM wavelengths, allowing them to travel simultaneously on the same fiber-optic cable. The CWDM OADM modules on the receiving end separate (or *demultiplex*) the different wavelengths.

For more information about the CWDM SFP modules and CWDM OADM modules, refer to the *Cisco CWDM GBIC and CWDM SFP Installation Note*.

Figure 1-4 Long-Distance, High-Bandwidth Transport Configuration



Where to Go Next

Before configuring the switch, review these sections for startup information:

- [Chapter 2, “Using the Command-Line Interface”](#)
- [Chapter 3, “Getting Started with CMS”](#)
- [Chapter 4, “Assigning the Switch IP Address and Default Gateway”](#)



Using the Command-Line Interface

This chapter describes the Cisco IOS command-line interface (CLI) and how to use it to configure your Catalyst 2970 switch. It contains these sections:

- [Understanding Command Modes, page 2-1](#)
- [Understanding the Help System, page 2-3](#)
- [Understanding Abbreviated Commands, page 2-4](#)
- [Understanding no and default Forms of Commands, page 2-4](#)
- [Understanding CLI Error Messages, page 2-5](#)
- [Using Command History, page 2-5](#)
- [Using Editing Features, page 2-6](#)
- [Searching and Filtering Output of show and more Commands, page 2-9](#)
- [Accessing the CLI, page 2-10](#)

Understanding Command Modes

The Cisco IOS user interface is divided into many different modes. The commands available to you depend on which mode you are currently in. Enter a question mark (?) at the system prompt to obtain a list of commands available for each command mode.

When you start a session on the switch, you begin in user mode, often called user EXEC mode. Only a limited subset of the commands are available in user EXEC mode. For example, most of the user EXEC commands are one-time commands, such as **show** commands, which show the current configuration status, and **clear** commands, which clear counters or interfaces. The user EXEC commands are not saved when the switch reboots.

To have access to all commands, you must enter privileged EXEC mode. Normally, you must enter a password to enter privileged EXEC mode. From this mode, you can enter any privileged EXEC command or enter global configuration mode.

Using the configuration modes (global, interface, and line), you can make changes to the running configuration. If you save the configuration, these commands are stored and used when the switch reboots. To access the various configuration modes, you must start at global configuration mode. From global configuration mode, you can enter interface configuration mode and line configuration mode.

Table 2-1 describes the main command modes, how to access each one, the prompt you see in that mode, and how to exit the mode. The examples in the table use the host name *Switch*.

Table 2-1 Command Mode Summary

Mode	Access Method	Prompt	Exit Method	About This Mode
User EXEC	Begin a session with your switch.	Switch>	Enter logout or quit .	Use this mode to <ul style="list-style-type: none"> • Change terminal settings. • Perform basic tests. • Display system information.
Privileged EXEC	While in user EXEC mode, enter the enable command.	Switch#	Enter disable to exit.	Use this mode to verify commands that you have entered. Use a password to protect access to this mode.
Global configuration	While in privileged EXEC mode, enter the configure command.	Switch(config)#	To exit to privileged EXEC mode, enter exit or end , or press Ctrl-Z .	Use this mode to configure parameters that apply to the entire switch.
Config-vlan	While in global configuration mode, enter the vlan <i>vlan-id</i> command.	Switch(config-vlan)#	To exit to global configuration mode, enter the exit command. To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure VLAN parameters. When VTP mode is transparent, you can create extended-range VLANs (VLAN IDs greater than 1005) and save configurations in the switch startup configuration file.
VLAN configuration	While in privileged EXEC mode, enter the vlan database command.	Switch(vlan)#	To exit to privileged EXEC mode, enter exit .	Use this mode to configure VLAN parameters for VLANs 1 to 1005 in the VLAN database.

Table 2-1 Command Mode Summary (continued)

Mode	Access Method	Prompt	Exit Method	About This Mode
Interface configuration	While in global configuration mode, enter the interface command (with a specific interface).	Switch(config-if)#	To exit to global configuration mode, enter exit . To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure parameters for the Ethernet interfaces. For information about defining interfaces, see the “Using Interface Configuration Mode” section on page 9-5. To configure multiple interfaces with the same parameters, see the “Configuring a Range of Interfaces” section on page 9-6.
Line configuration	While in global configuration mode, specify a line with the line vty or line console command.	Switch(config-line)#	To exit to global configuration mode, enter exit . To return to privileged EXEC mode, press Ctrl-Z or enter end .	Use this mode to configure parameters for the terminal line.

Understanding the Help System

You can enter a question mark (?) at the system prompt to display a list of commands available for each command mode. You can also obtain a list of associated keywords and arguments for any command, as shown in [Table 2-2](#).

Table 2-2 Help Summary

Command	Purpose
help	Obtain a brief description of the help system in any command mode.
<i>abbreviated-command-entry?</i>	Obtain a list of commands that begin with a particular character string. For example: <pre>Switch# di? dir disable disconnect</pre>
<i>abbreviated-command-entry<Tab></i>	Complete a partial command name. For example: <pre>Switch# sh conf<tab> Switch# show configuration</pre>
?	List all commands available for a particular command mode. For example: <pre>Switch> ?</pre>

Table 2-2 Help Summary (continued)

Command	Purpose
<i>command ?</i>	List the associated keywords for a command. For example: Switch> show ?
<i>command keyword ?</i>	List the associated arguments for a keyword. For example: Switch(config)# cdp holdtime ? <10-255> Length of time (in sec) that receiver must keep this packet

Understanding Abbreviated Commands

You need to enter only enough characters for the switch to recognize the command as unique.

This example shows how to enter the **show configuration** privileged EXEC command in an abbreviated form:

```
Switch# show conf
```

Understanding no and default Forms of Commands

Almost every configuration command also has a **no** form. In general, use the **no** form to disable a feature or function or reverse the action of a command. For example, the **no shutdown** interface configuration command reverses the shutdown of an interface. Use the command without the keyword **no** to re-enable a disabled feature or to enable a feature that is disabled by default.

Configuration commands can also have a **default** form. The **default** form of a command returns the command setting to its default. Most commands are disabled by default, so the **default** form is the same as the **no** form. However, some commands are enabled by default and have variables set to certain default values. In these cases, the **default** command enables the command and sets variables to their default values.

Understanding CLI Error Messages

Table 2-3 lists some error messages that you might encounter while using the CLI to configure your switch.

Table 2-3 Common CLI Error Messages

Error Message	Meaning	How to Get Help
% Ambiguous command: "show con"	You did not enter enough characters for your switch to recognize the command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command appear.
% Incomplete command.	You did not enter all the keywords or values required by this command.	Re-enter the command followed by a question mark (?) with a space between the command and the question mark. The possible keywords that you can enter with the command appear.
% Invalid input detected at '^' marker.	You entered the command incorrectly. The caret (^) marks the point of the error.	Enter a question mark (?) to display all the commands that are available in this command mode. The possible keywords that you can enter with the command appear.

Using Command History

The software provides a history or record of commands that you have entered. The command history feature is particularly useful for recalling long or complex commands or entries, including access lists. You can customize this feature to suit your needs as described in these sections:

- [Changing the Command History Buffer Size, page 2-5](#) (optional)
- [Recalling Commands, page 2-6](#) (optional)
- [Disabling the Command History Feature, page 2-6](#) (optional)

Changing the Command History Buffer Size

By default, the switch records ten command lines in its history buffer. You can alter this number for a current terminal session or for all sessions on a particular line. These procedures are optional.

Beginning in privileged EXEC mode, enter this command to change the number of command lines that the switch records during the current terminal session:

```
Switch# terminal history [size number-of-lines]
```

The range is from 0 to 256.

Beginning in line configuration mode, enter this command to configure the number of command lines the switch records for all sessions on a particular line:

```
Switch(config-line)# history [size number-of-lines]
```

The range is from 0 to 256.

Recalling Commands

To recall commands from the history buffer, perform one of the actions listed in [Table 2-4](#). These actions are optional.

Table 2-4 *Recalling Commands*

Action ¹	Result
Press Ctrl-P or the up arrow key.	Recall commands in the history buffer, beginning with the most recent command. Repeat the key sequence to recall successively older commands.
Press Ctrl-N or the down arrow key.	Return to more recent commands in the history buffer after recalling commands with Ctrl-P or the up arrow key. Repeat the key sequence to recall successively more recent commands.
show history	While in privileged EXEC mode, list the last several commands that you just entered. The number of commands that appear is determined by the setting of the terminal history global configuration command and history line configuration command.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Disabling the Command History Feature

The command history feature is automatically enabled. You can disable it for the current terminal session or for the command line. These procedures are optional.

To disable the feature during the current terminal session, enter the **terminal no history** privileged EXEC command.

To disable command history for the line, enter the **no history** line configuration command.

Using Editing Features

This section describes the editing features that can help you manipulate the command line. It contains these sections:

- [Enabling and Disabling Editing Features, page 2-7](#) (optional)
- [Editing Commands through Keystrokes, page 2-7](#) (optional)
- [Editing Command Lines that Wrap, page 2-8](#) (optional)

Enabling and Disabling Editing Features

Although enhanced editing mode is automatically enabled, you can disable it, re-enable it, or configure a specific line to have enhanced editing. These procedures are optional.

To globally disable enhanced editing mode, enter this command in line configuration mode:

```
Switch (config-line)# no editing
```

To re-enable the enhanced editing mode for the current terminal session, enter this command in privileged EXEC mode:

```
Switch# terminal editing
```

To reconfigure a specific line to have enhanced editing mode, enter this command in line configuration mode:

```
Switch(config-line)# editing
```

Editing Commands through Keystrokes

Table 2-5 shows the keystrokes that you need to edit command lines. These keystrokes are optional.

Table 2-5 Editing Commands through Keystrokes

Capability	Keystroke ¹	Purpose
Move around the command line to make changes or corrections.	Press Ctrl-B , or press the left arrow key.	Move the cursor back one character.
	Press Ctrl-F , or press the right arrow key.	Move the cursor forward one character.
	Press Ctrl-A .	Move the cursor to the beginning of the command line.
	Press Ctrl-E .	Move the cursor to the end of the command line.
	Press Esc B .	Move the cursor back one word.
	Press Esc F .	Move the cursor forward one word.
	Press Ctrl-T .	Transpose the character to the left of the cursor with the character located at the cursor.
Recall commands from the buffer and paste them in the command line. The switch provides a buffer with the last ten items that you deleted.	Press Ctrl-Y .	Recall the most recent entry in the buffer.
	Press Esc Y .	Recall the next buffer entry. The buffer contains only the last 10 items that you have deleted or cut. If you press Esc Y more than ten times, you cycle to the first buffer entry.
Delete entries if you make a mistake or change your mind.	Press the Delete or Backspace key.	Erase the character to the left of the cursor.

Table 2-5 Editing Commands through Keystrokes (continued)

Capability	Keystroke ¹	Purpose
	Press Ctrl-D .	Delete the character at the cursor.
	Press Ctrl-K .	Delete all characters from the cursor to the end of the command line.
	Press Ctrl-U or Ctrl-X .	Delete all characters from the cursor to the beginning of the command line.
	Press Ctrl-W .	Delete the word to the left of the cursor.
	Press Esc D .	Delete from the cursor to the end of the word.
Capitalize or lowercase words or capitalize a set of letters.	Press Esc C .	Capitalize at the cursor.
	Press Esc L .	Change the word at the cursor to lowercase.
	Press Esc U .	Capitalize letters from the cursor to the end of the word.
Designate a particular keystroke as an executable command, perhaps as a shortcut.	Press Ctrl-V or Esc Q .	
Scroll down a line or screen on displays that are longer than the terminal screen can display. Note The More prompt is used for any output that has more lines than can be displayed on the terminal screen, including show command output. You can use the Return and Space bar keystrokes whenever you see the More prompt.	Press the Return key.	Scroll down one line.
	Press the Space bar.	Scroll down one screen.
Redisplay the current command line if the switch suddenly sends a message to your screen.	Press Ctrl-L or Ctrl-R .	Redisplay the current command line.

1. The arrow keys function only on ANSI-compatible terminals such as VT100s.

Editing Command Lines that Wrap

You can use a wraparound feature for commands that extend beyond a single line on the screen. When the cursor reaches the right margin, the command line shifts ten spaces to the left. You cannot see the first ten characters of the line, but you can scroll back and check the syntax at the beginning of the command. The keystroke actions are optional.

To scroll back to the beginning of the command entry, press **Ctrl-B** or the left arrow key repeatedly. You can also press **Ctrl-A** to immediately move to the beginning of the line.



Note

The arrow keys function only on ANSI-compatible terminals such as VT100s.

In this example, the **access-list** global configuration command entry extends beyond one line. When the cursor first reaches the end of the line, the line is shifted ten spaces to the left and redisplayed. The dollar sign (\$) shows that the line has been scrolled to the left. Each time the cursor reaches the end of the line, the line is again shifted ten spaces to the left.

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1
Switch(config)# $ 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.25
Switch(config)# $t tcp 131.108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq
Switch(config)# $108.2.5 255.255.255.0 131.108.1.20 255.255.255.0 eq 45
```

After you complete the entry, press **Ctrl-A** to check the complete syntax before pressing the **Return** key to execute the command. The dollar sign (\$) appears at the end of the line to show that the line has been scrolled to the right:

```
Switch(config)# access-list 101 permit tcp 131.108.2.5 255.255.255.0 131.108.1$
```

The software assumes you have a terminal screen that is 80 columns wide. If you have a width other than that, use the **terminal width** privileged EXEC command to set the width of your terminal.

Use line wrapping with the command history feature to recall and modify previous complex command entries. For information about recalling previous command entries, see the [“Editing Commands through Keystrokes” section on page 2-7](#).

Searching and Filtering Output of show and more Commands

You can search and filter the output for **show** and **more** commands. This is useful when you need to sort through large amounts of output or if you want to exclude output that you do not need to see. Using these commands is optional.

To use this functionality, enter a **show** or **more** command followed by the *pipe* character (|), one of the keywords **begin**, **include**, or **exclude**, and an expression that you want to search for or filter out:

command | {**begin** | **include** | **exclude**} *regular-expression*

Expressions are case sensitive. For example, if you enter | **exclude output**, the lines that contain *output* are not displayed, but the lines that contain *Output* appear.

This example shows how to include in the output display only lines where the expression *protocol* appears:

```
Switch# show interfaces | include protocol
Vlan1 is up, line protocol is up
Vlan10 is up, line protocol is down
GigabitEthernet0/1 is up, line protocol is down
GigabitEthernet0/2 is up, line protocol is up
```

Accessing the CLI

You can access the CLI through a console connection, through Telnet, or by using the browser.

Accessing the CLI through a Console Connection or through Telnet

Before you can access the CLI, you must connect a terminal or PC to the switch console port and power on the switch as described in the hardware installation guide that shipped with your switch. Then, to understand the boot process and the options available for assigning IP information, see [Chapter 4, “Assigning the Switch IP Address and Default Gateway.”](#)

If your switch is already configured, you can access the CLI through a local console connection or through a remote Telnet session, but your switch must first be configured for this type of access. For more information, see the [“Setting a Telnet Password for a Terminal Line” section on page 7-6.](#)

You can use one of these methods to establish a connection with the switch:

- Connect the switch console port to a management station or dial-up modem. For information about connecting to the console port, refer to the switch hardware installation guide.
- Use any Telnet TCP/IP or encrypted Secure Shell (SSH) package from a remote management station. The switch must have network connectivity with the Telnet or SSH client, and the switch must have an enable secret password configured.

For information about configuring the switch for Telnet access, see the [“Setting a Telnet Password for a Terminal Line” section on page 7-6.](#) The switch supports up to 16 simultaneous Telnet sessions. Changes made by one Telnet user are reflected in all other Telnet sessions.

For information about configuring the switch for SSH, see the [“Configuring the Switch for Secure Shell” section on page 7-37.](#) The switch supports up to five simultaneous secure SSH sessions.

After you connect through the console port, through a Telnet session, or through an SSH session, the user EXEC prompt appears on the management station.

Accessing the CLI from a Browser

Before performing this procedure, make sure that you have met the software requirements (including browser and Java plug-in configurations) and have assigned IP information as described in the switch hardware installation guide. You also must assign a Telnet password to the switch or, if clustering, the command switch as described in [“Setting a Telnet Password for a Terminal Line” section on page 7-6.](#)

To access the CLI from a web browser, follow these steps:

-
- | | |
|--------|--|
| Step 1 | Start one of the supported browsers. |
| Step 2 | In the URL field, enter the IP address of the switch or, if clustering, the command switch. |
| Step 3 | When the Cisco Systems Access page appears, click Telnet to start a Telnet session. |
| Step 4 | Enter the switch password. |
- The user EXEC prompt appears on the management station.
-

**Note**

Copies of the HTML pages that you display are saved in your browser memory cache until you exit the browser session. A password is not required to redisplay these pages, including the Cisco Systems Access page. You can access the CLI by clicking **Web Console - HTML access to the command line interface** from a cached copy of the Cisco Systems Access page. To prevent unauthorized access to the CLI or to the Cluster Management Suite (CMS), exit your browser to end the browser session.



Getting Started with CMS

This chapter contains these sections that describe the Cluster Management Suite (CMS) on the Catalyst 2970 switch:

- “Understanding CMS” section on page 3-1
- “Configuring CMS” section on page 3-7
- “Displaying CMS” section on page 3-10
- “Where to Go Next” section on page 3-15

Refer to the appropriate switch documentation for descriptions of the browser-based management software used on other Catalyst switches. For more information about CMS, refer to the online help.

For a list of new CMS features in this release, select **Help > What’s New** from the CMS menu bar.

For information about cluster configurations and which Catalyst switches can be command switches or member switches, refer to the release notes for this switch.

Understanding CMS

CMS provides these features for managing switch clusters and individual switches from web browsers such as Netscape Communicator or Microsoft Internet Explorer:

- Front-panel and topology views of your network, as shown in [Figure 3-7 on page 3-13](#) and [Figure 3-8 on page 3-14](#), that can be displayed at the same time
- A menu bar, a toolbar, and a feature bar, as shown in [Figure 3-6 on page 3-13](#), to access configuration and management options
- Comprehensive online help that gives high-level concepts and procedures for performing CMS tasks
- Interactive modes—guide mode, expert mode, and wizards—that control the presentation of some complex configuration options
- Two levels of access modes to the configuration options: read-write access for users who can change switch settings and read-only access for users who can only view switch settings

Front Panel View

The Front Panel view displays the Front Panel image of a specific set of switches in a cluster. From this view, you can select multiple ports or multiple switches and configure them with the same settings.

For more information, see the [“Displaying CMS” section on page 3-10](#).

Topology View

The Topology view displays a network map that uses icons representing switch clusters, the command switch, cluster members, cluster candidates, neighboring devices that are not eligible to join a cluster, and link types. You can also display link information in the form of link reports and link graphs.

This view is available only when CMS is launched from a command switch.

For more information, see the [“Displaying CMS” section on page 3-10](#).

CMS Menu Bar, Toolbar, and Feature Bar

The configuration and monitoring options for configuring switches and switch clusters are available from the menu bar, the toolbar, and the feature bar.

- The menu bar, shown in [Figure 3-1](#), provides these options for managing a single switch and switch clusters:
 - CMS—Choose printing options, select interaction modes, display CMS preferences, save CMS cluster information to your PC or workstation, and show or hide the feature bar.

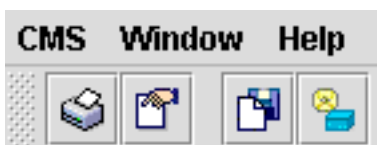


Note

CMS is downloaded to your browser each time that you launch CMS. You can increase the speed at which CMS loads by permanently installing CMS on your PC or workstation. Select **CMS > Installation and Distributions**, and click **Install**. CMS is installed locally and loads faster the next time you launch it.















- Window—Choose from the currently open CMS windows.
- Help—Launch the online help.

Figure 3-1 Menu Bar



- The toolbar provides buttons for commonly used switch and cluster configuration options and information windows such as legends and online help. [Table 3-1](#) lists the toolbar options from left to right on the toolbar.

Table 3-1 Toolbar Buttons

Toolbar Option	Icon	Task
Print		Print a CMS window or help file.
Preferences ¹		Set CMS display properties, such as polling intervals, the views to open at CMS startup, and the color of administratively shutdown ports.
Save Configuration ²		Save the configuration of the cluster or a switch to Flash memory.
Software Upgrade ²		Upgrade the software for the cluster or a switch.
Port Settings ¹		Display and configure port parameters on a switch.
VLAN ¹		Display VLAN membership, assign ports to VLANs, and change the administration mode.
Inventory		Display the device type, the software version, the IP address, and other information about a switch.
Refresh		Update the views with the latest status.
Front Panel		Display the Front Panel view.
Topology ³		Display the Topology view.
Topology Options ³		Select the information to be displayed in the Topology view.
Save Topology Layout ^{2 3}		Save your arrangement of the cluster icons in the Topology view to Flash memory.
Legend		Display the legend that describes the icons, labels, and links.
Help for Active Window		Display the help for the active, open window. You can also click Help from the active window or press the F1 key.

1. Not available in read-only mode. For more information about the read-only and read-write access modes, see the “[Privilege Levels](#)” section on page 3-6.

2. Some options from this menu option are not available in read-only mode.

3. Available only from a cluster-management session.

- The feature bar shows the features available for the devices in your cluster. By default, the feature bar is in standard mode. In this mode, the feature bar is always visible, and you can reduce or increase the width of the feature bar. In autohide mode, the feature bar appears only when you move the cursor to the left edge of the CMS workspace.
 - To enable the feature bar, click **CMS > Feature Bar**, and select **Standard Mode**.
 - To hide the feature bar, click **CMS > Feature Bar**, and select **Autohide Mode**.

Figure 3-2 shows the features available in a sample cluster.

Figure 3-2 Feature Bar and Search Window



1	Feature bar	2	Search window
---	-------------	---	---------------

**Note**

Only features supported by the devices in your cluster are displayed in the feature bar.

You can search for features that are available for your cluster by clicking **Search** and entering a feature name, as shown in [Figure 3-2](#).

Access modes affect the availability of features from CMS. Some CMS features are not available in read-only mode. For more information about how access modes affect CMS, see the [“Privilege Levels” section on page 3-6](#).

Online Help

CMS provides comprehensive online help to assist you in understanding and performing configuration and monitoring tasks from the CMS windows.

Online help is available for features that are supported by devices in your cluster. Sometimes the information in a topic differs for different cluster members. In these cases, the right pane contains all the versions of the topic, each labeled with the host names of the members it applies to.

Online help includes these features:

- Feature-specific help that gives background information and concepts on the features
- Dialog-specific help that gives procedures for performing tasks
- An index of online help topics
- A glossary of terms used in the online help

You can send us feedback about the information provided in the online help. Click **Feedback** to display an online form. After completing the form, click **Submit** to send your comments to Cisco Systems Inc. We appreciate and value your comments.

Configuration Modes

You can change the CMS interaction mode to either expert or guide mode. Expert mode displays a configuration window in which you configure the feature options. Guide mode takes you through each feature option and provides information about the parameter. Wizards are also available for some configuration options. These are similar to guide-mode configuration windows, except that fewer options are available.

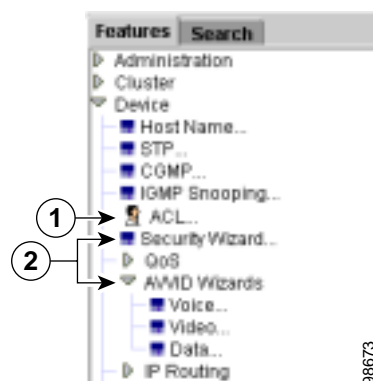
Guide Mode

Guide mode is for users who want a step-by-step approach for completing a specific configuration task. This mode is not available for all features. A person icon appears next to features that have guide mode available, as shown in [Figure 3-3](#).

When you click **Guide Mode** and then select a feature that supports it, CMS displays a specific parameter of that feature and information about the parameter. To configure the feature, you enter the information in each step until you click **Finish** in the last step. Clicking **Cancel** at any time ends the configuration task without applying any changes.

If you select **Guide Mode** but you want to use **Expert Mode** instead, you click **Guide** *before* selecting an option from the menu bar, tool bar, or popup menu. If you change the interaction mode after selecting a configuration option, the mode change does not take effect until you select another configuration option.

Figure 3-3 Guide Mode and Wizards



1	Guide mode icon	2	Wizards
---	-----------------	---	---------

Guide mode is not available if your switch access level is read-only. For more information about the read-only access mode, see the [“Privilege Levels” section on page 3-6](#).

Expert Mode

Expert mode is for users who prefer to display all the parameter fields of a feature in a single CMS window. You can view information about the parameter fields by clicking the **Help** button.

If you select **Expert Mode** but you want to use **Guide Mode** instead, you must click **Guide** *before* selecting an option from the menu bar, tool bar, or popup menu. If you change the interaction mode after selecting a configuration option, the mode change does not take effect until you select another configuration option.

Wizards

Similar to guide mode, wizards provide a step-by-step approach for completing a specific configuration task. Unlike guide mode, a wizard does not prompt you to provide information for all of the feature options. Instead, it prompts you to provide minimal information and then uses the default settings of the remaining options to set up default configurations.

When you select a feature that has *Wizard* in the name, the wizard launches for that feature, as shown in [Figure 3-3 on page 3-5](#).

Wizards are not available for all features or for read-only access levels. For more information about the read-only access mode, see the [“Privilege Levels” section on page 3-6](#).

Privilege Levels

CMS provides two levels of access to the configuration options: read-write access and read-only access. If you know your privilege level, you must specify it in the URL that you use to access the cluster. For example, if your privilege level is 13, enter this URL:

`http://ip_address/level/13`

Privilege levels 0 to 15 are supported.

- Privilege level 15 provides read-write access to CMS. This is the default.
- Privilege levels 1 to 14 provide read-only access to CMS. Any options in the CMS windows, menu bar, toolbar, and popup menus that change the switch or cluster configuration are not shown in read-only mode.
- Privilege level 0 denies access to CMS.

If you do not specify a privilege level when you access CMS, the switch verifies whether you have privilege level 15. If you do not, you are denied access to CMS. If you do have privilege level 15, you are granted read-write access. Therefore, you do not need to include the privilege level if it is 15. Entering zero denies access to CMS.

For more information about privilege levels, see the [“Preventing Unauthorized Access to Your Switch” section on page 7-1](#) and the [“Configuring Multiple Privilege Levels” section on page 7-8](#).

Access to Older Switches In a Cluster

If your cluster has these member switches running earlier software releases and if you have read-only access to these member switches, some configuration windows for those switches display incomplete information:

- Catalyst 2900 XL or Catalyst 3500 XL member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Catalyst 2950 member switches running Cisco IOS Release 12.0(5)WC2 or earlier

For more information about this limitation, refer to the release notes.

These switches do not support read-only mode on CMS:

- Catalyst 1900 and Catalyst 2820 switches
- Catalyst 2900 XL switches with 4-MB CPU DRAM

In read-only mode, these switches appear as unavailable devices and cannot be configured from CMS.

Configuring CMS

This section contains these topics that describe the requirements and configuration information for CMS:

- [“CMS Requirements” section on page 3-7](#)
- [“Cross-Platform Considerations” section on page 3-9](#)
- [“Launching CMS” section on page 3-10](#)

CMS Requirements

This section describes the hardware and software requirements for running CMS:

- [“Minimum Hardware Configuration” section on page 3-7](#)
- [“Operating System and Browser Support” section on page 3-8](#)
- [“Browser Plug-In Requirements” section on page 3-8](#)
- [“Specifying an HTTP Port \(Nondefault Configuration Only\)” section on page 3-9](#)
- [“Configuring an Authentication Method \(Nondefault Configuration Only\)” section on page 3-10](#)

**Note**

The software requirements are automatically verified by the CMS Startup Report when you launch CMS. For more information, see the [“Launching CMS” section on page 3-10](#).

Minimum Hardware Configuration

The minimum PC requirement is a Pentium processor running at 233 MHz with 64 MB of DRAM. The minimum UNIX workstation requirement is a Sun Ultra 1 running at 143 MHz with 64 MB of DRAM.

[Table 3-2](#) lists the minimum platforms for running CMS.

Table 3-2 Minimum Hardware Configuration

OS	Processor Speed	DRAM	Number of Colors	Resolution	Font Size
Windows NT 4.0 ¹	Pentium 300 MHz	128 MB	65,536	1024 x 768	Small
Solaris 2.5.1 or higher	SPARC 333 MHz	128 MB	Most colors for applications	—	Small (3)

1. Service Pack 3 or higher is required.

Operating System and Browser Support

You can access the CMS interface by using the operating systems and browsers listed in [Table 3-3](#). CMS checks the browser version when starting a session to ensure that the browser is supported.

Table 3-3 Supported Operating Systems and Browsers

Operating System	Minimum Service Pack or Patch	Netscape Communicator ¹	Microsoft Internet Explorer ²
Windows 98	Second Edition	7.1	5.5 or 6.0
Windows NT 4.0	Service Pack 3 or later	7.1	5.5 or 6.0
Windows 2000	None	7.1	5.5 or 6.0
Windows XP	None	7.1	5.5 or 6.0
Solaris 2.5.1 or later	Sun-recommended patch cluster for the OS and Motif library patch 103461-24	7.0	Not supported

1. Netscape Communicator version 6.0 is not supported.

2. Service Pack 1 or higher is required for Internet Explorer 5.5.

Browser Plug-In Requirements

You need to install a browser plug-in to run CMS.

Windows

For Windows platforms, the CMS plug-in is required to run CMS. For more information about the CMS plug-in, including the URL, see the “Software Compatibility” section in the release notes.



Note

If you need to both upgrade your web browser and install the CMS plug-in, you *must* upgrade your browser first. If you install the CMS plug-in and then upgrade your browser, the plug-in is not registered with the new browser.



Note

Do not install the CMS plug-in on Solaris.

Solaris

For Solaris, Java plug-in 1.4.1 is required to run CMS. You can download the Java plug-in and installation instructions from this URL:

<http://www.cisco.com/cgi-bin/tablebuild.pl/java>

On Solaris platforms, follow the instructions in the README_FIRST.txt file to install the Java plug-in.

You need to close and restart your browser after installing a Java plug-in.

Cross-Platform Considerations

When managing switch clusters through CMS, remember that clusters can have a mix of switch models using different Cisco IOS releases and that CMS in earlier Cisco IOS releases and on different switch platforms might look and function differently from CMS in this Cisco IOS release.

When you select **Device > Device Manager** for a cluster member, a new browser session launches, and the CMS version for that switch appears.

Here are examples of how CMS can differ between Cisco IOS releases and switch platforms:

- On Catalyst switches running Cisco IOS Release 12.0(5)WC2 or earlier or Cisco IOS Release 12.1(6)EA1 or earlier, the CMS versions in those software releases might appear similar but are not the same as this release. For example, the Topology view in this release is not the same as the Topology view or the Cluster View in those earlier software releases.
- CMS on the Catalyst 1900 and Catalyst 2820 switches is referred to as Switch Manager. Cluster management options are not available on these switches. This is the earliest version of CMS.

Refer to the documentation specific to the switch and its Cisco IOS release for descriptions of the CMS version.

HTTP Access to CMS

CMS uses the HTTP protocol (the default is port 80) and the default method of authentication (the enable password) to communicate with the switch through any of its Ethernet ports and to allow switch management from a standard web browser.

If you have not configured a specific (nondefault) HTTP port and are using the enable password (or no password) for access to the switch, you can go to the “[Displaying CMS](#)” section on page 3-10.

Specifying an HTTP Port (Nondefault Configuration Only)

If you change the HTTP port, you must include the new port number when you enter the IP address in the browser **Location** or **Address** field (for example, <http://10.1.126.45:184> where 184 is the new HTTP port number.) You should write down the port number to which you are connected. Use care when changing the switch IP information.

Configuring an Authentication Method (Nondefault Configuration Only)

If you are *not* using the default method of authentication (the enable password), you need to configure the HTTP server interface with the method of authentication used on the switch.

Beginning in privileged EXEC mode, follow these steps to configure the HTTP server interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip http authentication { enable local tacacs }	Configure the HTTP server interface for the type of authentication you want to use. <ul style="list-style-type: none"> • enable—Enable password, which is the default method of HTTP server user authentication. • local—Local user database as defined on the Cisco router or access server is used. • tacacs—TACACS server is used.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.

After you have configured the HTTP server interface, display the CMS access page, as described in the [“Launching CMS” section on page 3-10](#).

Displaying CMS

This section provides these topics about displaying CMS:

[“Launching CMS” section on page 3-10](#)

[“Front Panel View” section on page 3-13](#)

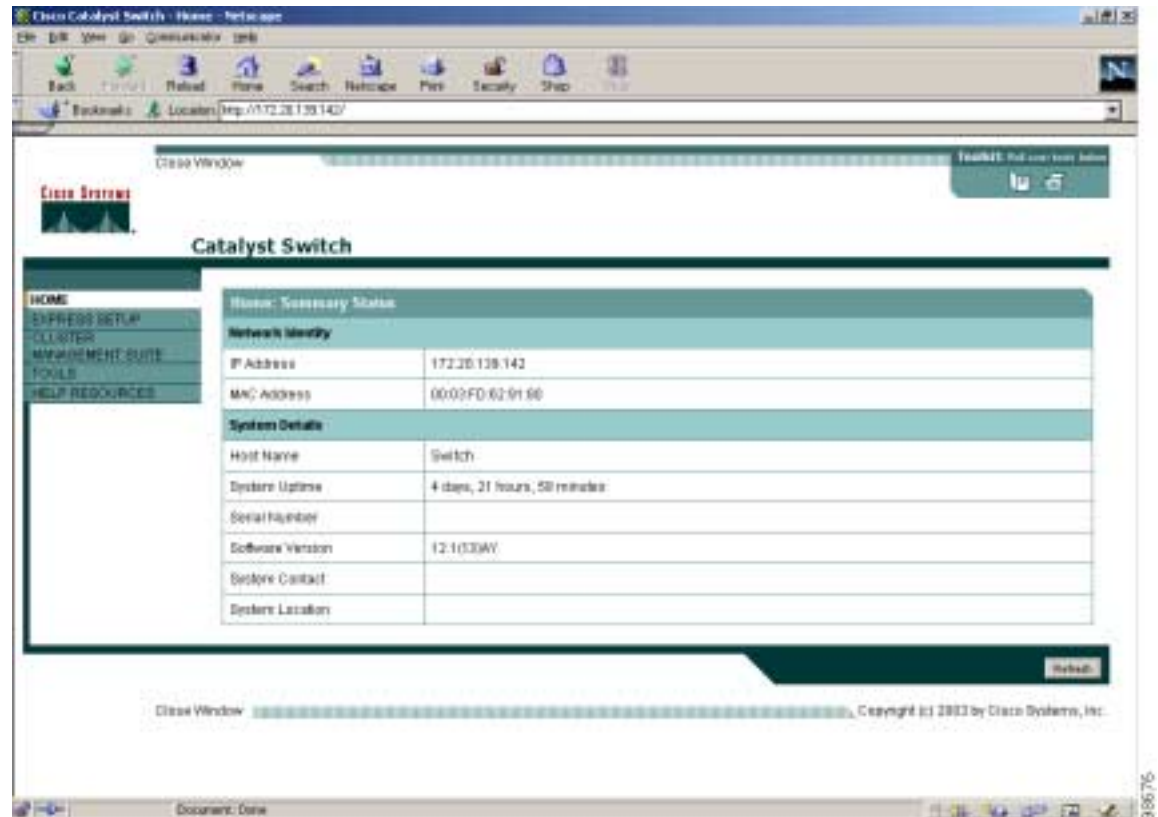
[“Topology View” section on page 3-14](#)

Launching CMS

To display the switch access page, follow these steps:

-
- | | |
|--------|--|
| Step 1 | Enter the switch IP address in the browser, and press Return . |
| Step 2 | Enter your username and password when prompted. If no username is configured on your switch (the default), enter only the enable password (if an enable password is configured) in the password field. |
- The switch home page appears, as shown in [Figure 3-4](#).

Figure 3-4 Switch Home Page



The Switch Home Page has these tabs:

- Express Setup—Opens the Express Setup page



Note You can use Express Setup to assign an IP address to an unconfigured switch. For more information, refer to the hardware installation guide.

- Cluster Management Suite—Launches CMS
- Tools—Accesses diagnostic and monitoring tools, such as Telnet, Extended Ping, and the **show interfaces** privileged EXEC command
- Help Resources—Provides links to the Cisco website, technical documentation, and the Cisco Technical Assistance Center (TAC)

Step 3 Click **Cluster Management Suite** to launch the CMS interface. The CMS Startup Report runs and verifies that your PC or workstation can correctly run CMS.

If you are running an unsupported operating system, web browser, CMS plug-in or Java plug-in, or if the plug-in is not enabled, the CMS Startup Report page appears, as shown in [Figure 3-5](#).

Figure 3-5 CMS Startup Report



The CMS Startup Report has links that instruct you how to correctly configure your PC or workstation. If the CMS Startup Report appears, click the links, and follow the instructions to configure your PC or workstation.



Note

If you are running Windows and need to both upgrade your web browser and install the CMS plug-in, you *must* upgrade your browser first. If you install the CMS plug-in and then upgrade your browser, the plug-in is not registered with the new browser.



Note

If your PC or workstation is correctly configured for CMS, you do not see the CMS Startup Report.

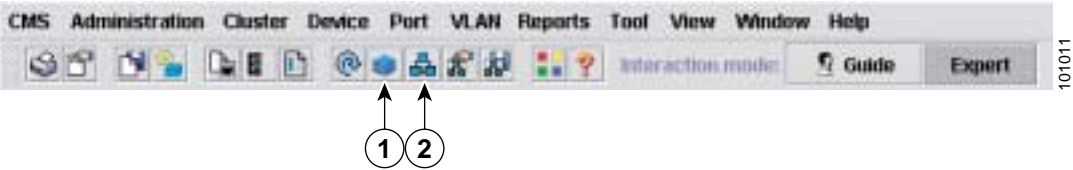
When your PC or workstation is correctly configured, CMS launches.

Front Panel View

When CMS is launched from a noncommand switch, the Front Panel view displays by default, and the front-panel view displays only the front panel of the specific switch.

When CMS is launched from a command switch, you can display the Front Panel view by clicking the Front Panel button on the tool bar, as shown in [Figure 3-6](#).

Figure 3-6 Toolbar

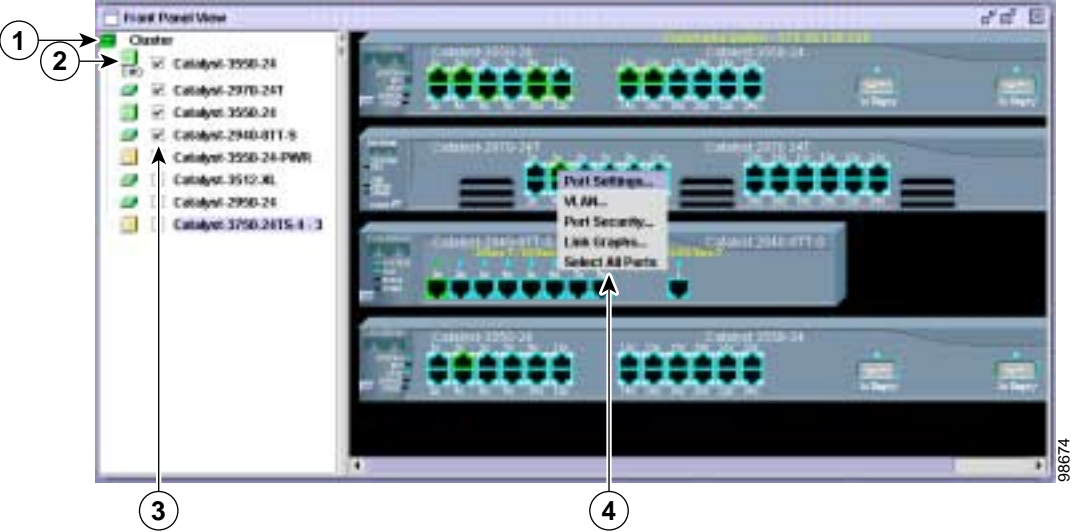


1	Front Panel view button	2	Topology view button
---	-------------------------	---	----------------------

The Front Panel view displays the front-panel image of the command switch and other selected switches, as shown in [Figure 3-7](#), and you can select more switches to be displayed.

You can choose and configure the switches that appear in Front Panel view. You can drag the switches that appear and re-arrange them. You can right-click on a switch port to configure that port.

Figure 3-7 Front Panel View and Port Popup Menu



1	Cluster tree	3	Checkboxes to show switches
2	Command switch	4	Port configuration popup menu

**Note**

Figure 3-7 shows a cluster with a Catalyst 3550 switch as the command switch. Refer to the release notes for a list of switches that can be members of a cluster with a Catalyst 2970 switch as the command switch.

**Note**

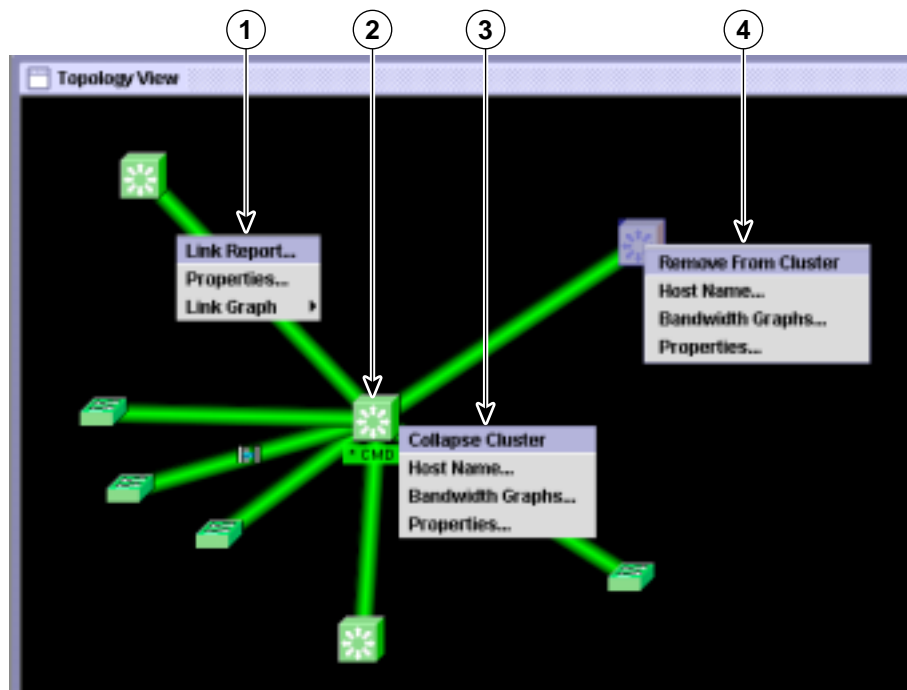
On Catalyst 1900 and Catalyst 2820 switches, CMS is referred to as Device Manager (also referred to as Switch Manager). Device Manager is for configuring an individual switch. When you select Device Manager for a specific switch in the cluster, you launch a separate CMS session. The Device Manager interface can vary among the Catalyst switch platforms.

Topology View

When CMS is launched from a command switch, the Topology view appears by default. (This view is available only when CMS is launched from a command switch.)

When you click the topology button on the tool bar, the Topology view displays the command switch (indicated by the **CMD** label) and the devices that are connected to it, as shown in Figure 3-8. You can right-click on a switch or link icon to display a menu for that icon.

Figure 3-8 Topology View and Device Popup Menus



1	Link popup menu	3	Command switch popup menu
2	Command switch	4	Cluster member popup menu

**Note**

Figure 3-8 shows multiple popup menus. Only one popup menu at a time appears in the CMS.

The Topology view shows how the devices within a switch cluster are connected and how the switch cluster is connected to other clusters and devices. From this view, you can add and remove cluster members. This view provides two levels of detail of the network topology:

- **Expand Cluster**—When you right-click a cluster icon and select **Expand Cluster**, the Topology view displays the switch cluster in detail. This view shows the command switch and member switches in a cluster. It also shows candidate switches that can join the cluster. This view does not display the details of any neighboring switch clusters
- **Collapse Cluster**—When you right-click a command-switch icon and select **Collapse Cluster**, the cluster is collapsed and represented by a single icon. The view shows how the cluster is connected to other clusters, candidate switches, and devices that are not eligible to join the cluster (such as routers, access points, IP phones, and so on).

**Note**

The Topology view displays only the switch cluster and network neighborhood of the specific command or member switch that you access. To display a different switch cluster, you need to access the command switch or member switch of that cluster.

CMS Icons

For a complete list of device and link icons available in CMS, select **Help > Legend** from the CMS menu bar.

Where to Go Next

- See [Chapter 5, “Clustering Switches,”](#) for more information about command and member switches.
- See [Chapter 6, “Administering the Switch,”](#) for more information about administrative tasks.
- Click **Help > What’s New** in the online help for a list of new CMS features in this release.

The rest of this guide provides information about the command-line interface (CLI) procedures for the software features supported in this release. For CMS procedures and window descriptions, refer to the online help.



Assigning the Switch IP Address and Default Gateway

This chapter describes how to create the initial switch configuration (for example, assigning the switch IP address and default gateway information) for the Catalyst 2970 switch by using a variety of automatic and manual methods. It also describes how to modify the switch startup configuration.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding the Boot Process, page 4-1](#)
- [Assigning Switch Information, page 4-2](#)
- [Checking and Saving the Running Configuration, page 4-11](#)
- [Modifying the Startup Configuration, page 4-12](#)
- [Scheduling a Reload of the Software Image, page 4-17](#)

Understanding the Boot Process

To start your switch, you need to follow the procedures in the hardware installation guide about installing and powering on the switch, and setting up the initial configuration (IP address, subnet mask, default gateway, secret and Telnet passwords, and so forth) of the switch.

The normal boot process involves the operation of the boot loader software, which performs these activities:

- Performs low-level CPU initialization. It initializes the CPU registers, which control where physical memory is mapped, its quantity, its speed, and so forth.
- Performs power-on self-test (POST) for the CPU subsystem. It tests the CPU DRAM and the portion of the Flash device that makes up the Flash file system.
- Initializes the Flash file system on the system board.
- Loads a default operating system software image into memory and boots the switch.

The boot loader provides access to the Flash file system before the operating system is loaded. Normally, the boot loader is used only to load, uncompress, and launch the operating system. After the boot loader gives the operating system control of the CPU, the boot loader is not active until the next system reset or power-on.

The boot loader also provides trap-door access into the system if the operating system has problems serious enough that it cannot be used. The trap-door mechanism provides enough access to the system so that if it is necessary, you can format the Flash file system, reinstall the operating system software image by using the XMODEM Protocol, recover from a lost or forgotten password, and finally restart the operating system. For more information, see the [“Recovering from Corrupted Software By Using the XMODEM Protocol” section on page 29-2](#) and the [“Recovering from a Lost or Forgotten Password” section on page 29-4](#).



Note

You can disable password recovery. For more information, see the [“Disabling Password Recovery” section on page 7-5](#).

Before you can assign switch information, make sure you have connected a PC or terminal to the console port, and configured the PC or terminal-emulation software baud rate and character format to match these of the switch console port:

- Baud rate default is 9600.
- Data bits default is 8.



Note

If the data bits option is set to 8, set the parity option to none.

- Stop bits default is 1.
- Parity settings default is none.

Assigning Switch Information

You can assign IP information through the switch setup program, through a Dynamic Host Configuration Protocol (DHCP) server, or manually.

Use the switch setup program if you are a new user and want to be prompted for specific IP information. With this program, you can also configure a host name and an enable secret password. It gives you the option of assigning a Telnet password (to provide security during remote management) and configuring your switch as a command or member switch of a cluster or as a standalone switch. For more information about the setup program, refer to the release notes on Cisco.com.

Use a DHCP server for centralized control and automatic assignment of IP information once the server is configured.

**Note**

If you are using DHCP, do not respond to any of the questions in the setup program until the switch receives the dynamically assigned IP address and reads the configuration file.

Use the manual method of configuration if you are an experienced user familiar with the switch configuration steps; otherwise, use the setup program described earlier.

This section contains this configuration information:

- [Default Switch Information, page 4-3](#)
- [Understanding DHCP-Based Autoconfiguration, page 4-3](#)
- [Manually Assigning IP Information, page 4-10](#)

Default Switch Information

[Table 4-1](#) shows the default switch information.

Table 4-1 *Default Switch Information*

Feature	Default Setting
IP address and subnet mask	No IP address or subnet mask are defined.
Default gateway	No default gateway is defined.
Enable secret password	No password is defined.
Host name	The factory-assigned default host name is <i>Switch</i> .
Telnet password	No password is defined.
Cluster command switch functionality	Disabled.
Cluster name	No cluster name is defined.

Understanding DHCP-Based Autoconfiguration

The DHCP provides configuration information to Internet hosts and internetworking devices. This protocol consists of two components: one for delivering configuration parameters from a DHCP server to a device and a mechanism for allocating network addresses to devices. DHCP is built on a client-server model, in which designated DHCP servers allocate network addresses and deliver configuration parameters to dynamically configured devices. The switch can act as both a DHCP client and a DHCP server.

During DHCP-based autoconfiguration, your switch (DHCP client) is automatically configured at startup with IP address information and a configuration file.

With DHCP-based autoconfiguration, no DHCP client-side configuration is needed on your switch. However, you need to configure the DHCP server, or the DHCP server feature on your switch, for various lease options associated with IP addresses. If you are using DHCP to relay the configuration file location on the network, you might also need to configure a Trivial File Transfer Protocol (TFTP) server and a Domain Name System (DNS) server.

The DHCP server, or the DHCP server feature running on your switch, can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. A relay device forwards broadcast traffic between two directly connected LANs. A router does not forward broadcast packets, but it forwards packets based on the destination IP address in the received packet.

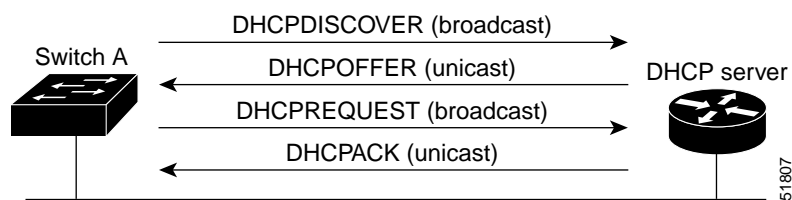
DHCP-based autoconfiguration replaces the BOOTP client functionality on your switch.

DHCP Client Request Process

When you boot your switch, the DHCP client is invoked and automatically requests configuration information from a DHCP server when the configuration file is not present on the switch.

[Figure 4-1](#) shows the sequence of messages that are exchanged between the DHCP client and the DHCP server.

Figure 4-1 DHCP Client and Server Message Exchange



The client, Switch A, broadcasts a DHCPDISCOVER message to locate a DHCP server. The DHCP server offers configuration parameters (such as an IP address, subnet mask, gateway IP address, DNS IP address, a lease for the IP address, and so forth) to the client in a DHCPOFFER unicast message.

In a DHCPREQUEST broadcast message, the client returns a formal request for the offered configuration information to the DHCP server. The formal request is broadcast so that all other DHCP servers that received the DHCPDISCOVER broadcast message from the client can reclaim the IP addresses that they offered to the client.

The DHCP server confirms that the IP address has been allocated to the client by returning a DHCPACK unicast message to the client. With this message, the client and server are bound, and the client uses configuration information received from the server. The amount of information the switch receives depends on how you configure the DHCP server. For more information, see the [“Configuring the DHCP Server”](#) section on page 4-5.

If the configuration parameters sent to the client in the DHCPOFFER unicast message are invalid (a configuration error exists), the client returns a DHCPDECLINE broadcast message to the DHCP server.

The DHCP server sends the client a DHCPNAK denial broadcast message, which means that the offered configuration parameters have not been assigned, that an error has occurred during the negotiation of the parameters, or that the client has been slow in responding to the DHCPOFFER message (the DHCP server assigned the parameters to another client).

A DHCP client might receive offers from multiple DHCP or BOOTP servers and can accept any of the offers; however, the client usually accepts the first offer it receives. The offer from the DHCP server is not a guarantee that the IP address is allocated to the client; however, the server usually reserves the address until the client has had a chance to formally request the address. If the switch accepts replies from a BOOTP server and configures itself, the switch broadcasts, instead of unicasts, TFTP requests to obtain the switch configuration file.

Configuring DHCP-Based Autoconfiguration

These sections describe how to configure DHCP-based autoconfiguration.

- [Configuring the DHCP Server, page 4-5](#)
- [Configuring the TFTP Server, page 4-6](#)
- [Configuring the DNS, page 4-6](#)
- [Configuring the Relay Device, page 4-6](#)
- [Obtaining Configuration Files, page 4-7](#)
- [Example Configuration, page 4-8](#)

If your DHCP server is a Cisco device, or if you are configuring the switch as a DHCP server, refer to the “*IP Addressing and Services*” section in the *Cisco IOS IP and IP Routing Configuration Guide for Cisco IOS Release 12.1* for additional information about configuring DHCP.

Configuring the DHCP Server

The switch can act as both the DHCP client and DHCP server. By default, the Cisco IOS DHCP server and relay agent features are enabled on your switch.

You should configure the DHCP server or the DHCP server feature running on your switch with reserved leases that are bound to each switch by the switch hardware address.

If you want the switch to receive IP address information, you must configure the DHCP server with these lease options:

- IP address of the client (required)
- Subnet mask of the client (required)
- DNS server IP address (optional)
- Router IP address (default gateway address to be used by the switch) (required)

If you want the switch to receive the configuration file from a TFTP server, you must configure the DHCP server with these lease options:

- TFTP server name (required)
- Boot filename (the name of the configuration file that the client needs) (recommended)
- Host name (optional)

Depending on the settings of the DHCP server, the switch can receive IP address information, the configuration file, or both.

If you do not configure the DHCP server or the DHCP server feature running on your switch with the lease options described earlier, it replies to client requests with only those parameters that are configured. If the IP address and subnet mask are not in the reply, the switch is not configured. If the router IP address or TFTP server name are not found, the switch might send broadcast, instead of unicast, TFTP requests. Unavailability of other lease options does not affect autoconfiguration.

The DHCP server or the DHCP server feature running on your switch can be on the same LAN or on a different LAN than the switch. If the DHCP server is running on a different LAN, you should configure a DHCP relay. For more information, see the “[Configuring the Relay Device](#)” section on page 4-6.

Configuring the TFTP Server

Based on the DHCP server configuration, the switch attempts to download one or more configuration files from the TFTP server. If you configured the DHCP server to respond to the switch with all the options required for IP connectivity to the TFTP server, and if you configured the DHCP server with a TFTP server name, address, and configuration filename, the switch attempts to download the specified configuration file from the specified TFTP server.

If you did not specify the configuration filename, the TFTP server, or if the configuration file could not be downloaded, the switch attempts to download a configuration file by using various combinations of filenames and TFTP server addresses. The files include the specified configuration filename (if any) and these files: `network-config`, `cisconet.cfg`, `hostname.config`, or `hostname.cfg`, where *hostname* is the switch's current hostname. The TFTP server addresses used include the specified TFTP server address (if any) and the broadcast address (255.255.255.255).

For the switch to successfully download a configuration file, the TFTP server must contain one or more configuration files in its base directory. The files can include these files:

- The configuration file named in the DHCP reply (the actual switch configuration file).
- The `network-config` or the `cisconet.cfg` file (known as the default configuration files).
- The `router-config` or the `ciscotr.cfg` file (These files contain commands common to all switches. Normally, if the DHCP and TFTP servers are properly configured, these files are not accessed.)

If you specify the TFTP server name in the DHCP server-lease database, you must also configure the TFTP server name-to-IP-address mapping in the DNS-server database.

If the TFTP server to be used is on a different LAN from the switch, or if it is to be accessed by the switch through the broadcast address (which occurs if the DHCP server response does not contain all the required information described earlier), a relay must be configured to forward the TFTP packets to the TFTP server. For more information, see the [“Configuring the Relay Device” section on page 4-6](#). The preferred solution is to configure the DHCP server or the DHCP server feature running on your switch with all the required information.

Configuring the DNS

The DHCP server or the DHCP server feature running on your switch uses the DNS server to resolve the TFTP server name to an IP address. You must configure the TFTP server name-to-IP address map on the DNS server. The TFTP server contains the configuration files for the switch.

You can configure the IP addresses of the DNS servers in the lease database of the DHCP server from where the DHCP replies will retrieve them. You can enter up to two DNS server IP addresses in the lease database.

The DNS server can be on the same or on a different LAN as the switch. If it is on a different LAN, the switch must be able to access it through a router.

Configuring the Relay Device

You must configure a relay device when a switch sends broadcast packets that need to be responded to by a host on a different LAN. Examples of broadcast packets that the switch might send are DHCP, DNS, and in some cases, TFTP packets. You must configure this relay device to forward received broadcast packets on an interface to the destination host.

If the relay device is a Cisco router, enable IP routing (**ip routing** global configuration command), and configure helper addresses by using the **ip helper-address** interface configuration command.

For example, in Figure 4-2, configure the router interfaces as follows:

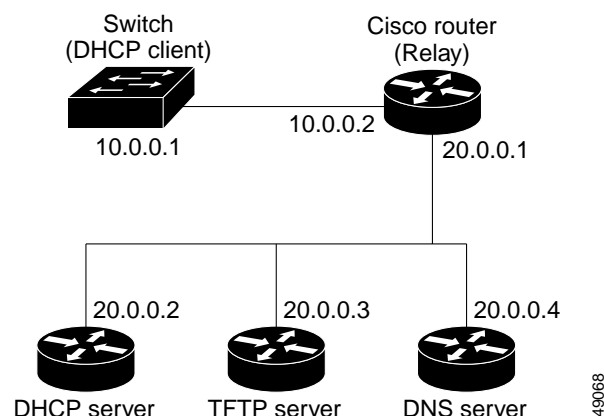
On interface 10.0.0.2:

```
router(config-if)# ip helper-address 20.0.0.2
router(config-if)# ip helper-address 20.0.0.3
router(config-if)# ip helper-address 20.0.0.4
```

On interface 20.0.0.1

```
router(config-if)# ip helper-address 10.0.0.1
```

Figure 4-2 Relay Device Used in Autoconfiguration



Obtaining Configuration Files

Depending on the availability of the IP address and the configuration filename in the DHCP reserved lease, the switch obtains its configuration information in these ways:

- The IP address and the configuration filename is reserved for the switch and provided in the DHCP reply (one-file read method).

The switch receives its IP address, subnet mask, TFTP server address, and the configuration filename from the DHCP server or the DHCP server feature running on your switch. The switch sends a unicast message to the TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- The IP address and the configuration filename is reserved for the switch, but the TFTP server address is not provided in the DHCP reply (one-file read method).

The switch receives its IP address, subnet mask, and the configuration filename from the DHCP server, or the DHCP server feature running on your switch. The switch sends a broadcast message to a TFTP server to retrieve the named configuration file from the base directory of the server, and upon receipt, completes its boot-up process.

- Only the IP address is reserved for the switch and provided in the DHCP reply. The configuration filename is not provided (two-file read method).

The switch receives its IP address, subnet mask, and the TFTP server address from the DHCP server or the DHCP server feature running on your switch. The switch sends a unicast message to the TFTP server to retrieve the network-config or cisco.net.cfg default configuration file. (If the network-config file cannot be read, the switch reads the cisco.net.cfg file.)

The default configuration file contains the host names-to-IP-address mapping for the switch. The switch fills its host table with the information in the file and obtains its host name. If the host name is not found in the file, the switch uses the host name in the DHCP reply. If the host name is not specified in the DHCP reply, the switch uses the default *Switch* as its host name.

After obtaining its host name from the default configuration file or the DHCP reply, the switch reads the configuration file that has the same name as its host name (*hostname-conf* or *hostname.cfg*, depending on whether *network-conf* or *cisconet.cfg* was read earlier) from the TFTP server. If the *cisconet.cfg* file is read, the filename of the host is truncated to eight characters.

If the switch cannot read the *network-conf*, *cisconet.cfg*, or the *hostname* file, it reads the *router-conf* file. If the switch cannot read the *router-conf* file, it reads the *ciscotr.cfg* file.

**Note**

The switch broadcasts TFTP server requests if the TFTP server is not obtained from the DHCP replies, if all attempts to read the configuration file through unicast transmissions fail, or if the TFTP server name cannot be resolved to an IP address.

Example Configuration

Figure 4-3 shows a sample network for retrieving IP information by using DHCP-based autoconfiguration.

Figure 4-3 DHCP-Based Autoconfiguration Network Example

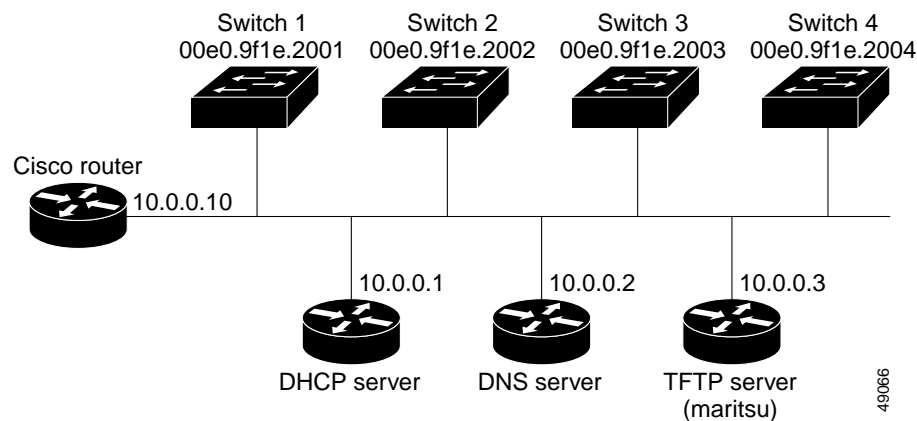


Table 4-2 shows the configuration of the reserved leases on the DHCP server or the DHCP server feature running on your switch.

Table 4-2 DHCP Server Configuration

	Switch-1	Switch-2	Switch-3	Switch-4
Binding key (hardware address)	00e0.9f1e.2001	00e0.9f1e.2002	00e0.9f1e.2003	00e0.9f1e.2004
IP address	10.0.0.21	10.0.0.22	10.0.0.23	10.0.0.24
Subnet mask	255.255.255.0	255.255.255.0	255.255.255.0	255.255.255.0
Router address	10.0.0.10	10.0.0.10	10.0.0.10	10.0.0.10
DNS server address	10.0.0.2	10.0.0.2	10.0.0.2	10.0.0.2
TFTP server name	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3	maritsu or 10.0.0.3
Boot filename (configuration file) (optional)	switch1-config	switch2-config	switch3-config	switch4-config
Host name (optional)	switch1	switch2	switch3	switch4

DNS Server Configuration

The DNS server maps the TFTP server name *maritsu* to IP address 10.0.0.3.

TFTP Server Configuration (on UNIX)

The TFTP server base directory is set to `/tftpserver/work/`. This directory contains the `network-config` file used in the two-file read method. This file contains the host name to be assigned to the switch based on its IP address. The base directory also contains a configuration file for each switch (*switch1-config*, *switch2-config*, and so forth) as shown in this display:

```
prompt> cd /tftpserver/work/
prompt> ls
network-config
switch1-config
switch2-config
switch3-config
switch4-config
prompt> cat network-config
ip host switch1 10.0.0.21
ip host switch2 10.0.0.22
ip host switch3 10.0.0.23
ip host switch4 10.0.0.24
```

DHCP Client Configuration

No configuration file is present on Switch 1 through Switch 4.

Configuration Explanation

In Figure 4-3, Switch 1 reads its configuration file as follows:

- It obtains its IP address 10.0.0.21 from the DHCP server.
- If no configuration filename is given in the DHCP server reply, Switch 1 reads the `network-config` file from the base directory of the TFTP server.
- It adds the contents of the `network-config` file to its host table.
- It reads its host table by indexing its IP address 10.0.0.21 to its host name (switch1).

- It reads the configuration file that corresponds to its host name; for example, it reads *switch1-conf* from the TFTP server.

Switches 2 through 4 retrieve their configuration files and IP addresses in the same way.

Manually Assigning IP Information

Beginning in privileged EXEC mode, follow these steps to manually assign IP information to multiple switched virtual interfaces (SVIs) or ports:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface vlan <i>vlan-id</i>	Enter interface configuration mode, and enter the VLAN to which the IP information is assigned. The range is 1 to 4094; do not enter leading zeros.
Step 3	ip address <i>ip-address subnet-mask</i>	Enter the IP address and subnet mask.
Step 4	exit	Return to global configuration mode.
Step 5	ip default-gateway <i>ip-address</i>	<p>Enter the IP address of the next-hop router interface that is directly connected to the switch where a default gateway is being configured. The default gateway receives IP packets with unresolved destination IP addresses from the switch.</p> <p>Once the default gateway is configured, the switch has connectivity to the remote networks with which a host needs to communicate.</p> <p>Note When your switch is configured to route with IP, it does not need to have a default gateway set.</p>
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the switch IP address, use the **no ip address** interface configuration command. If you are removing the address through a Telnet session, your connection to the switch will be lost. To remove the default gateway address, use the **no ip default-gateway** global configuration command.

For information on setting the switch system name, protecting access to privileged EXEC commands, and setting time and calendar services, see [Chapter 6, “Administering the Switch.”](#)

Checking and Saving the Running Configuration

You can check the configuration settings you entered or changes you made by entering this privileged EXEC command:

```
Switch# show running-config
Building configuration...

Current configuration: 1363 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname Switch A
!
enable secret 5 $l$ej9.$DMUvAUUnZOAmvmgqBEzIxE0
!
interface gigabitethernet0/1
 no switchport
 ip address 172.20.137.50 255.255.255.0
!
interface gigabitethernet0/2
!
interface gigabitethernet0/3
 mvr type source
!
interface gigabitethernet0/4
!
interface gigabitethernet0/5
!
interface gigabitethernet0/6
!
interface gigabitethernet0/7
!
interface gigabitethernet0/8
!
interface gigabitethernet0/9
 no ip address
!
interface gigabitethernet0/10
!
interface gigabitethernet0/11
!
interface gigabitethernet0/12
...!
interface VLAN1
 ip address 172.20.137.50 255.255.255.0
 no ip directed-broadcast
!
ip default-gateway 172.20.137.1 !
!
snmp-server community private RW
snmp-server community public RO
snmp-server community private@es0 RW
snmp-server community public@es0 RO
snmp-server chassis-id 0x12
!
end
```

To store the configuration or changes you have made to your startup configuration in Flash memory, enter this privileged EXEC command:

```
Switch# copy running-config startup-config
Destination filename [startup-config]?
Building configuration...
```

This command saves the configuration settings that you made. If you fail to do this, your configuration will be lost the next time you reload the system. To display information stored in the NVRAM section of Flash memory, use the **show startup-config** or **more startup-config** privileged EXEC command.

For more information about alternative locations to copy the configuration file, see [Appendix B, “Working with the Cisco IOS File System, Configuration Files, and Software Images.”](#)

Modifying the Startup Configuration

This section describes how to modify the switch startup configuration. It contains this configuration information:

- [Default Boot Configuration, page 4-12](#)
- [Automatically Downloading a Configuration File, page 4-13](#)
- [Booting Manually, page 4-13](#)
- [Booting a Specific Software Image, page 4-14](#)
- [Controlling Environment Variables, page 4-15](#)

See also [Appendix B, “Working with the Cisco IOS File System, Configuration Files, and Software Images,”](#) for information about switch configuration files.

Default Boot Configuration

[Table 4-3](#) shows the default boot configuration.

Table 4-3 Default Boot Configuration

Feature	Default Setting
Operating system software image	<p>The switch attempts to automatically boot the system using information in the BOOT environment variable. If the variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the Flash file system.</p> <p>The Cisco IOS image is stored in a directory that has the same name as the image file (excluding the .bin extension).</p> <p>In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory.</p>
Configuration file	<p>Configured switches use the <i>config.text</i> file stored on the system board in Flash memory.</p> <p>A new switch has no configuration file.</p>

Automatically Downloading a Configuration File

You can automatically download a configuration file to your switch by using the DHCP-based autoconfiguration feature. For more information, see the [“Understanding DHCP-Based Autoconfiguration” section on page 4-3](#).

Specifying the Filename to Read and Write the System Configuration

By default, the Cisco IOS software uses the file *config.text* to read and write a nonvolatile copy of the system configuration. However, you can specify a different filename, which will be loaded during the next boot cycle.

Beginning in privileged EXEC mode, follow these steps to specify a different configuration filename:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot config-file flash:/file-url	Specify the configuration file to load during the next boot cycle. For <i>file-url</i> , specify the path (directory) and the configuration filename. Filenames and directory names are case sensitive.
Step 3	end	Return to privileged EXEC mode.
Step 4	show boot	Verify your entries. The boot config-file global configuration command changes the setting of the CONFIG_FILE environment variable.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot config-file** global configuration command.

Booting Manually

By default, the switch automatically boots; however, you can configure it to manually boot.

Beginning in privileged EXEC mode, follow these steps to configure the switch to manually boot during the next boot cycle:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot manual	Enable the switch to manually boot during the next boot cycle.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show boot	<p>Verify your entries.</p> <p>The boot manual global command changes the setting of the MANUAL_BOOT environment variable.</p> <p>The next time you reboot the system, the switch is in boot loader mode, shown by the <i>switch:</i> prompt. To boot the system, use the boot filesystem:/file-url boot loader command.</p> <ul style="list-style-type: none"> For <i>filesystem:</i>, use flash: for the system board Flash device. For <i>file-url</i>, specify the path (directory) and the name of the bootable image. <p>Filenames and directory names are case sensitive.</p>
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable manual booting, use the **no boot manual** global configuration command.

Booting a Specific Software Image

By default, the switch attempts to automatically boot the system using information in the BOOT environment variable. If this variable is not set, the switch attempts to load and execute the first executable image it can by performing a recursive, depth-first search throughout the Flash file system. In a depth-first search of a directory, each encountered subdirectory is completely searched before continuing the search in the original directory. However, you can specify a specific image to boot.

Beginning in privileged EXEC mode, follow these steps to configure the switch to boot a specific image during the next boot cycle:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	boot system filesystem:/file-url	<p>Configure the switch to boot a specific image in Flash memory during the next boot cycle.</p> <ul style="list-style-type: none"> For <i>filesystem:</i>, use flash: for the system board Flash device. For <i>file-url</i>, specify the path (directory) and the name of the bootable image. <p>Filenames and directory names are case sensitive.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show boot	<p>Verify your entries.</p> <p>The boot system global command changes the setting of the BOOT environment variable.</p> <p>During the next boot cycle, the switch attempts to automatically boot the system using information in the BOOT environment variable.</p>
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no boot system** global configuration command.

Controlling Environment Variables

With a normally operating switch, you enter the boot loader mode only through a switch console connection configured for 9600 bps. Unplug the switch power cord, and press the switch **Mode** button while reconnecting the power cord. You can release the **Mode** button a second or two after the LED above port 1 turns off. Then the boot loader *switch:* prompt appears.

The switch boot loader software provides support for nonvolatile environment variables, which can be used to control how the boot loader, or any other software running on the system, behaves. Boot loader environment variables are similar to environment variables that can be set on UNIX or DOS systems.

Environment variables that have values are stored in Flash memory outside of the Flash file system.

Each line in these files contains an environment variable name and an equal sign followed by the value of the variable. A variable has no value if it is not listed in this file; it has a value if it is listed in the file even if the value is a null string. A variable that is set to a null string (for example, “”) is a variable with a value. Many environment variables are predefined and have default values.

Environment variables store two kinds of data:

- Data that controls code, which does not read the Cisco IOS configuration file. For example, the name of a boot loader helper file, which extends or patches the functionality of the boot loader can be stored as an environment variable.
- Data that controls code, which is responsible for reading the Cisco IOS configuration file. For example, the name of the Cisco IOS configuration file can be stored as an environment variable.

You can change the settings of the environment variables by accessing the boot loader or by using Cisco IOS commands. Under normal circumstances, it is not necessary to alter the setting of the environment variables.



Note

For complete syntax and usage information for the boot loader commands and environment variables, refer to the command reference for this release.

Table 4-4 describes the function of the most common environment variables.

Table 4-4 Environment Variables

Variable	Boot Loader Command	Cisco IOS Global Configuration Command
BOOT	<p>set BOOT <i>filesystem:/file-url ...</i></p> <p>A semicolon-separated list of executable files to try to load and execute when automatically booting. If the BOOT environment variable is not set, the system attempts to load and execute the first executable image it can find by using a recursive, depth-first search through the Flash file system. If the BOOT variable is set but the specified images cannot be loaded, the system attempts to boot the first bootable file that it can find in the Flash file system.</p>	<p>boot system <i>filesystem:/file-url</i></p> <p>Specifies the Cisco IOS image to load during the next boot cycle. This command changes the setting of the BOOT environment variable.</p>
MANUAL_BOOT	<p>set MANUAL_BOOT <i>yes</i></p> <p>Determines whether the switch automatically or manually boots.</p> <p>Valid values are 1, yes, 0, and no. If it is set to no or 0, the boot loader attempts to automatically boot the system. If it is set to anything else, you must manually boot the switch from the boot loader mode.</p>	<p>boot manual</p> <p>Enables manually booting the switch during the next boot cycle and changes the setting of the MANUAL_BOOT environment variable.</p> <p>The next time you reboot the system, the switch is in boot loader mode. To boot the system, use the boot loader boot flash:<i>filesystem:/file-url</i> command, and specify the name of the bootable image.</p>
CONFIG_BUFSIZE	<p>set CONFIG_BUFSIZE <i>size</i></p> <p>Changes the buffer size that Cisco IOS uses to hold a copy of the configuration file in memory. The configuration file cannot be larger than the buffer size allocation. The range is from 4096 to 524288 bytes.</p>	<p>boot buffersize <i>size</i></p> <p>Specifies the size of the file system-simulated NVRAM in Flash memory. The buffer holds a copy of the configuration file in memory. This command changes the setting of the CONFIG_BUFSIZE environment variable.</p> <p>You must reload the switch by using the reload privileged EXEC command for this command to take effect.</p>
CONFIG_FILE	<p>set CONFIG_FILE <i>flash:/file-url</i></p> <p>Changes the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration.</p>	<p>boot config-file <i>flash:/file-url</i></p> <p>Specifies the filename that Cisco IOS uses to read and write a nonvolatile copy of the system configuration. This command changes the CONFIG_FILE environment variable.</p>

Scheduling a Reload of the Software Image

You can schedule a reload of the software image to occur on the switch at a later time (for example, late at night or during the weekend when the switch is used less), or you can synchronize a reload network-wide (for example, to perform a software upgrade on all switches in the network).



Note

A scheduled reload must take place within approximately 24 days.

Configuring a Scheduled Reload

To configure your switch to reload the software image at a later time, use one of these commands in privileged EXEC mode:

- **reload in** *[hh:]mm* *[text]*

This command schedules a reload of the software to take affect in the specified minutes or hours and minutes. The reload must take place within approximately 24 days. You can specify the reason for the reload in a string up to 255 characters in length.

- **reload at** *hh:mm* *[month day | day month]* *[text]*

This command schedules a reload of the software to take place at the specified time (using a 24-hour clock). If you specify the month and day, the reload is scheduled to take place at the specified time and date. If you do not specify the month and day, the reload takes place at the specified time on the current day (if the specified time is later than the current time) or on the next day (if the specified time is earlier than the current time). Specifying 00:00 schedules the reload for midnight.



Note

Use the **at** keyword only if the switch system clock has been set (through Network Time Protocol (NTP), the hardware calendar, or manually). The time is relative to the configured time zone on the switch. To schedule reloads across several switches to occur simultaneously, the time on each switch must be synchronized with NTP.

The **reload** command halts the system. If the system is not set to manually boot, it reboots itself. Use the **reload** command after you save the switch configuration information to the startup configuration (**copy running-config startup-config**).

If your switch is configured for manual booting, do not reload it from a virtual terminal. This restriction prevents the switch from entering the boot loader mode and thereby taking it from the remote user's control.

If you modify your configuration file, the switch prompts you to save the configuration before reloading. During the save operation, the system requests whether you want to proceed with the save if the CONFIG_FILE environment variable points to a startup configuration file that no longer exists. If you proceed in this situation, the system enters setup mode upon reload.

This example shows how to reload the software on the switch on the current day at 7:30 p.m:

```
Switch# reload at 19:30
Reload scheduled for 19:30:00 UTC Wed Jun 5 1996 (in 2 hours and 25 minutes)
Proceed with reload? [confirm]
```

This example shows how to reload the software on the switch at a future time:

```
Switch# reload at 02:00 jun 20
Reload scheduled for 02:00:00 UTC Thu Jun 20 1996 (in 344 hours and 53 minutes)
Proceed with reload? [confirm]
```

To cancel a previously scheduled reload, use the **reload cancel** privileged EXEC command.

Displaying Scheduled Reload Information

To display information about a previously scheduled reload or to determine if a reload has been scheduled on the switch, use the **show reload** privileged EXEC command.

It displays reload information including the time the reload is scheduled to occur and the reason for the reload (if it was specified when the reload was scheduled).



Clustering Switches

This chapter provides the concepts and procedures to create and manage Catalyst 2970 switch clusters.



Note

This chapter focuses on Catalyst 2970 switch clusters. It also includes guidelines and limitations for clusters mixed with other cluster-capable Catalyst switches, but it does not provide complete descriptions of the cluster features for these other switches. For complete cluster information for a specific Catalyst platform, refer to the software configuration guide for that switch.

This chapter consists of these sections:

- [Understanding Switch Clusters, page 5-2](#)
- [Planning a Switch Cluster, page 5-4](#)
- [Creating a Switch Cluster, page 5-15](#)



Note

Configuring switch clusters is more easily done from the Cluster Management Suite (CMS) web-based interface than through the command-line interface (CLI). Therefore, information in this chapter focuses on using CMS to create a cluster. See [Chapter 3, “Getting Started with CMS,”](#) for additional information about switch clusters and the clustering options. For complete procedures about using CMS to configure switch clusters, refer to the online help. For the CLI cluster commands, refer to the switch command reference.

- [Verifying a Switch Cluster, page 5-20](#)
- [Using the CLI to Manage Switch Clusters, page 5-21](#)
- [Using SNMP to Manage Switch Clusters, page 5-21](#)

Understanding Switch Clusters

A *switch cluster* is a set of up to 16 connected, cluster-capable Catalyst switches that are managed as a single entity. The switches in the cluster use the switch clustering technology so that you can configure and troubleshoot a group of different Catalyst desktop switch platforms through a single IP address.

In a switch cluster, 1 switch must be the *cluster command switch* and up to 15 other switches can be *cluster member switches*. The total number of switches in a cluster cannot exceed 16 switches. The cluster command switch is the single point of access used to configure, manage, and monitor the cluster member switches. Cluster members can belong to only one cluster at a time.

The benefits of clustering switches include:

- Management of Catalyst switches regardless of their interconnection media and their physical locations. The switches can be in the same location, or they can be distributed across a Layer 2 or Layer 3 (if your cluster is using a Catalyst 3550 or Catalyst 3750 switch as a Layer 3 router between the Layer 2 switches in the cluster) network.

Cluster members are connected to the cluster command switch according to the connectivity guidelines described in the [“Automatic Discovery of Cluster Candidates and Members” section on page 5-5](#). This section includes management VLAN considerations for the Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

- Command-switch redundancy if a cluster command switch fails. One or more switches can be designated as *standby cluster command switches* to avoid loss of contact with cluster members. A *cluster standby group* is a group of standby cluster command switches.
- Management of a variety of Catalyst switches through a single IP address. This conserves on IP addresses, especially if you have a limited number of them. All communication with the switch cluster is through the cluster command switch IP address.

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and the required software versions.

These sections describe:

- [Cluster Command Switch Characteristics, page 5-3](#)
- [Standby Cluster Command Switch Characteristics, page 5-3](#)
- [Candidate Switch and Cluster Member Switch Characteristics, page 5-4](#)

Cluster Command Switch Characteristics

A cluster command switch must meet these requirements:

- It is running Cisco IOS Release 12.1(11)AX or later.
- It has an IP address.
- It has Cisco Discovery Protocol (CDP) version 2 enabled (the default).
- It is not a command or cluster member switch of another cluster.
- It is connected to the standby cluster command switches through the management VLAN and to the cluster member switches through a common VLAN.

**Note**

If your switch cluster has a Catalyst 2970 switch, it should be the cluster command switch unless the cluster has a Catalyst 3750 switch or switch stack. If the switch cluster has a Catalyst 3750 switch or switch stack, that switch or switch stack must be the cluster command switch.

Standby Cluster Command Switch Characteristics

A standby cluster command switch must meet these requirements:

- It is running Cisco IOS Release 12.1(11)AX or later.
- It has an IP address.
- It has CDP version 2 enabled.
- It is connected to the command switch and to other standby command switches through its management VLAN.
- It is connected to all other cluster member switches (except the cluster command and standby command switches) through a common VLAN.
- It is redundantly connected to the cluster so that connectivity to cluster member switches is maintained.
- It is not a command or member switch of another cluster.

**Note**

Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2970 switch, the standby cluster command switches must also be Catalyst 2970 switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

Candidate Switch and Cluster Member Switch Characteristics

Candidate switches are cluster-capable switches that have not yet been added to a cluster. Cluster member switches are switches that have actually been added to a switch cluster. Although not required, a candidate or cluster member switch can have its own IP address and password (for related considerations, see the [“IP Addresses” section on page 5-13](#) and [“Passwords” section on page 5-14](#)).

To join a cluster, a candidate switch must meet these requirements:

- It is running cluster-capable software.
- It has CDP version 2 enabled.
- It is not a command or cluster member switch of another cluster.
- If a cluster standby group exists, it is connected to every standby cluster command switch through at least one common VLAN. The VLAN to each standby cluster command switch can be different.
- It is connected to the cluster command switch through at least one common VLAN.

**Note**

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL candidate and cluster member switches must be connected through their management VLAN to the cluster command switch and standby cluster command switches. For complete information about these switches in a switch-cluster environment, refer to the software configuration guide for that specific switch.

This requirement does not apply if you have a Catalyst 2970, Catalyst 3550, or Catalyst 3750 cluster command switch. Candidate and cluster member switches can connect through any VLAN in common with the cluster command switch.

Planning a Switch Cluster

Anticipating conflicts and compatibility issues is a high priority when you manage several switches through a cluster. This section describes these guidelines, requirements, and caveats that you should understand before you create the cluster:

- [Automatic Discovery of Cluster Candidates and Members, page 5-5](#)
- [HSRP and Standby Cluster Command Switches, page 5-10](#)
- [IP Addresses, page 5-13](#)
- [Host Names, page 5-13](#)
- [Passwords, page 5-14](#)
- [SNMP Community Strings, page 5-14](#)
- [TACACS+ and RADIUS, page 5-14](#)
- [Access Modes in CMS, page 5-15](#)
- [LRE Profiles, page 5-15](#)
- [Availability of Switch-Specific Features in Switch Clusters, page 5-15](#)

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and for the required software versions and browser and Java plug-in configurations.

Automatic Discovery of Cluster Candidates and Members

The cluster command switch uses Cisco Discovery Protocol (CDP) to discover cluster member switches, candidate switches, neighboring switch clusters, and edge devices across multiple VLANs and in star or cascaded topologies.

**Note**

Do not disable CDP on the cluster command switch, on cluster members, or on any cluster-capable switches that you might want a cluster command switch to discover. For more information about CDP, see [Chapter 20, “Configuring CDP.”](#)

Following these connectivity guidelines ensures automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices:

- [Discovery Through CDP Hops, page 5-5](#)
- [Discovery Through Non-CDP-Capable and Noncluster-Capable Devices, page 5-6](#)
- [Discovery Through Different VLANs, page 5-7](#)
- [Discovery Through Different Management VLANs, page 5-8](#)
- [Discovery of Newly Installed Switches, page 5-9](#)

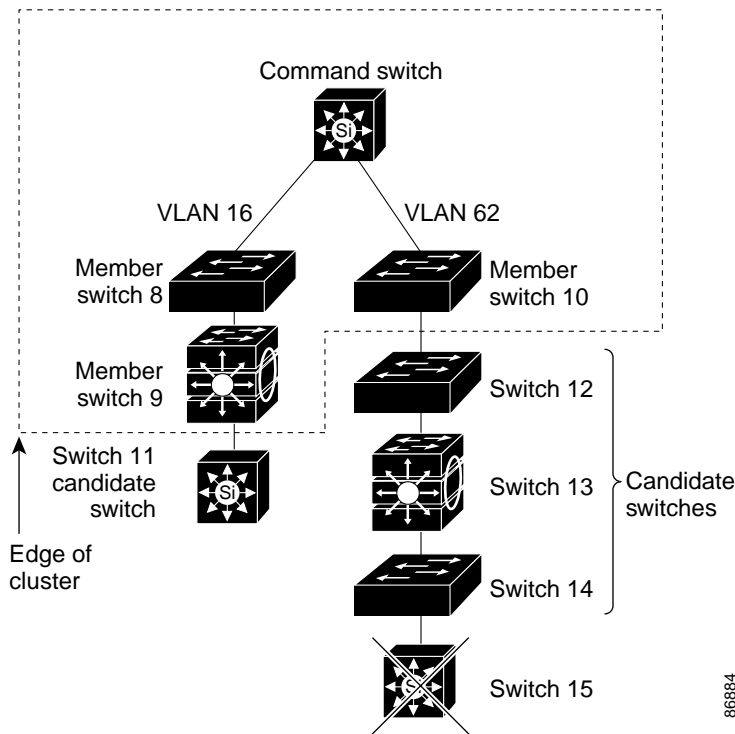
Discovery Through CDP Hops

By using CDP, a cluster command switch can discover switches up to seven CDP hops away (the default is three hops) from the edge of the cluster. The edge of the cluster is where the last cluster member switches are connected to the cluster and to candidate switches. For example, cluster member switches 9 and 10 in [Figure 5-1](#) are at the edge of the cluster.

You can set the number of hops the cluster command switch searches for candidate and cluster member switches by selecting **Cluster > Hop Count**. When new candidate switches are added to the network, the cluster command switch discovers them and adds them to the list of candidate switches.

In [Figure 5-1](#), the cluster command switch has ports assigned to VLANs 16 and 62. The CDP hop count is three. The cluster command switch discovers switches 11, 12, 13, and 14 because they are within three hops from the edge of the cluster. It does not discover switch 15 because it is four hops from the edge of the cluster.

Figure 5-1 Discovery Through CDP Hops

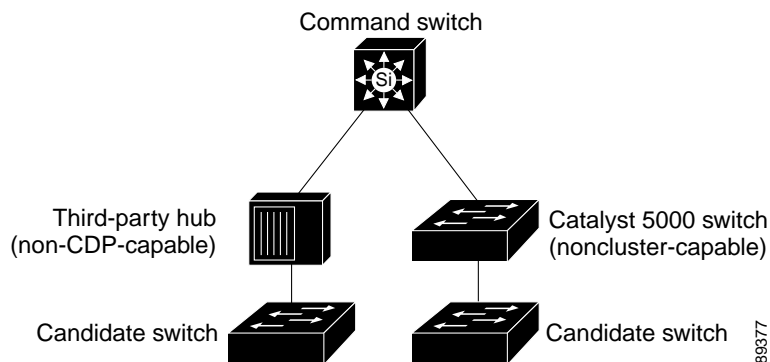


Discovery Through Non-CDP-Capable and Noncluster-Capable Devices

If a cluster command switch is connected to a *non-CDP-capable third-party hub* (such as a non-Cisco hub), it can discover cluster-enabled devices connected to that third-party hub. However, if the cluster command switch is connected to a *noncluster-capable Cisco device*, it cannot discover a cluster-enabled device connected beyond the noncluster-capable Cisco device.

Figure 5-2 shows that the cluster command switch discovers the switch that is connected to a third-party hub. However, the cluster command switch does not discover the switch that is connected to a Catalyst 5000 switch.

Figure 5-2 Discovery Through Non-CDP-Capable and Noncluster-Capable Devices

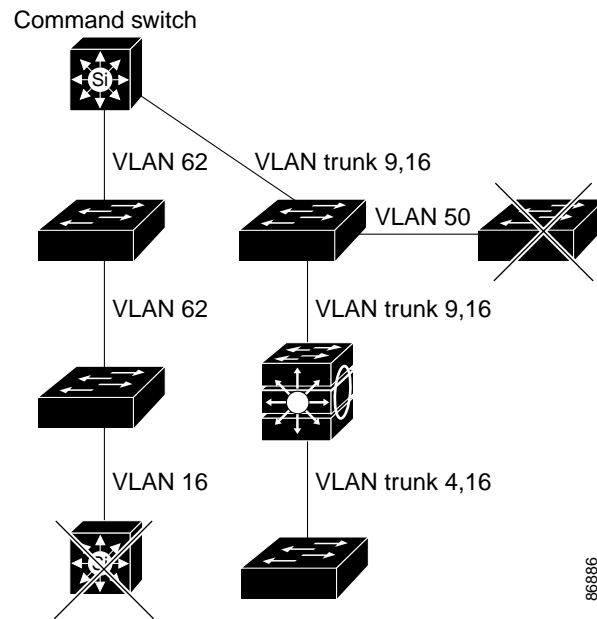


Discovery Through Different VLANs

If the cluster command switch is a Catalyst 2970, Catalyst 3550, or Catalyst 3750 switch, the cluster can have cluster member switches in different VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. The cluster command switch in [Figure 5-3](#) has ports assigned to VLANs 9, 16, and 62 and therefore discovers the switches in those VLANs. It does not discover the switch in VLAN 50. It also does not discover the switch in VLAN 16 in the first column because the cluster command switch has no VLAN connectivity to it.

Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster command switch through their management VLAN. For information about discovery through management VLANs, the [“Discovery Through Different Management VLANs”](#) section on page 5-8. For more information about VLANs, see [Chapter 11, “Configuring VLANs.”](#)

Figure 5-3 Discovery Through Different VLANs



Discovery Through Different Management VLANs

Catalyst 2970, Catalyst 3550, or Catalyst 3750 cluster command switches can discover and manage cluster member switches in different VLANs and different management VLANs. As cluster member switches, they must be connected through at least one VLAN in common with the cluster command switch. They do not need to be connected to the cluster command switch through their management VLAN. The default management VLAN is VLAN 1.



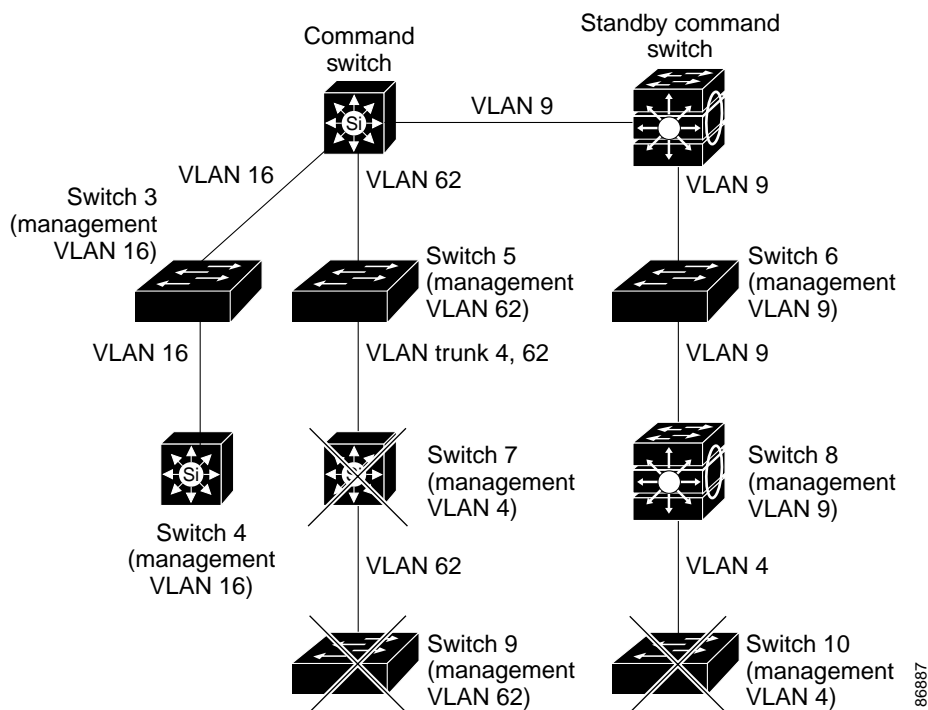
Note

If the switch cluster has a Catalyst 3750 switch or switch stack, that switch or switch stack must be the cluster command switch.

The cluster command switch and standby command switch in [Figure 5-4](#) have ports assigned to VLANs 9, 16, and 62. The management VLAN on the cluster command switch is VLAN 9. Each cluster command switch discovers the switches in the different management VLANs except these:

- Switches 7 and 10 (switches in management VLAN 4) because they are not connected through a common VLAN (meaning VLANs 62 and 9) with the cluster command switch
- Switch 9 because automatic discovery does not extend beyond a noncandidate device, which is switch 7

Figure 5-4 Discovery Through Different Management VLANs with a Layer 3 Cluster Command Switch



Discovery of Newly Installed Switches

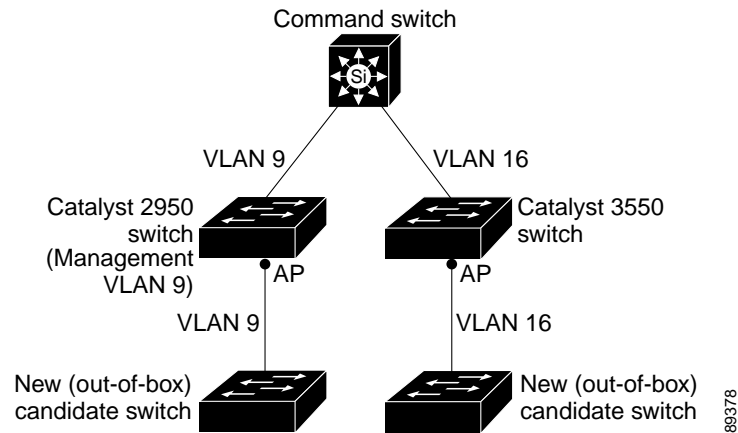
To join a cluster, the new, out-of-the-box switch must be connected to the cluster through one of its access ports. An access port (AP) carries the traffic of and belongs to only one VLAN. By default, the new switch and its access ports are assigned to VLAN 1.

When the new switch joins a cluster, its default VLAN changes to the VLAN of the immediately upstream neighbor. The new switch also configures its access port to belong to the VLAN of the immediately upstream neighbor.

The cluster command switch in [Figure 5-5](#) belongs to VLANs 9 and 16. When new cluster-capable switches join the cluster:

- One cluster-capable switch and its access port are assigned to VLAN 9.
- The other cluster-capable switch and its access port are assigned to management VLAN 16.

Figure 5-5 Discovery of Newly Installed Switches



HSRP and Standby Cluster Command Switches

The switch uses Hot Standby Router Protocol (HSRP) so that you can configure a group of standby cluster command switches. Because a cluster command switch manages the forwarding of all communication and configuration information to all the cluster member switches, we strongly recommend the following:

- For a cluster command switch stack, a standby cluster command switch is necessary if the entire switch stack fails. However, if only the stack master in the command switch stack fails, the switch stack elects a new stack master and resumes its role as the cluster command switch stack.
- For a cluster command switch that is a standalone switch, configure a standby cluster command switch to take over if the primary cluster command switch fails.

A *cluster standby group* is a group of command-capable switches that meet the requirements described in the “[Standby Cluster Command Switch Characteristics](#)” section on page 5-3. Only one cluster standby group can be assigned per cluster.



Note

If your switch cluster has a Catalyst 2970 switch, it should be the cluster command switch unless the cluster has a Catalyst 3750 switch or switch stack. If the switch cluster has a Catalyst 3750 switch or switch stack, that switch or switch stack must be the cluster command switch.

The switches in the cluster standby group are ranked according to HSRP priorities. The switch with the highest priority in the group is the *active cluster command switch* (AC). The switch with the next highest priority is the *standby cluster command switch* (SC). The other switches in the cluster standby group are the *passive cluster command switches* (PC). If the active cluster command switch and the standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. For the limitations to automatic discovery, see the “[Automatic Recovery of Cluster Configuration](#)” section on page 5-12. For information about changing HSRP priority values, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*. The HSRP **standby priority** interface configuration commands are the same for changing the priority of cluster standby group members and router-redundancy group members.



Note

The HSRP standby hold time interval should be greater than or equal to three times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and standby hello time intervals, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

These connectivity guidelines ensure automatic discovery of the switch cluster, cluster candidates, connected switch clusters, and neighboring edge devices. These topics also provide more detail about standby cluster command switches:

- [Virtual IP Addresses, page 5-11](#)
- [Other Considerations for Cluster Standby Groups, page 5-11](#)
- [Automatic Recovery of Cluster Configuration, page 5-12](#)

Virtual IP Addresses

You need to assign a unique virtual IP address and group number and name to the cluster standby group. This information must be configured on a specific VLAN or routed port on the active cluster command switch. The active cluster command switch receives traffic destined for the virtual IP address. To manage the cluster, you must access the active cluster command switch through the virtual IP address, not through the command-switch IP address. This is in case the IP address of the active cluster command switch is different from the virtual IP address of the cluster standby group.

If the active cluster command switch fails, the standby cluster command switch assumes ownership of the virtual IP address and becomes the active cluster command switch. The passive switches in the cluster standby group compare their assigned priorities to determine the new standby cluster command switch. The passive standby switch with the highest priority then becomes the standby cluster command switch. When the previously active cluster command switch becomes active again, it resumes its role as the active cluster command switch, and the current active cluster command switch becomes the standby cluster command switch again. For more information about IP address in switch clusters, see the [“IP Addresses” section on page 5-13](#).

Other Considerations for Cluster Standby Groups

These requirements also apply:

- Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2970 switch, the standby cluster command switches must also be Catalyst 2970 switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

If your switch cluster has a Catalyst 2970 switch, it should be the cluster command switch unless the cluster has a Catalyst 3750 switch or switch stack. If the switch cluster has a Catalyst 3750 switch or switch stack, that switch or switch stack must be the cluster command switch.

- Only one cluster standby group can be assigned to a cluster. You can have more than one router-redundancy standby group.
- All standby-group members must be members of the cluster.



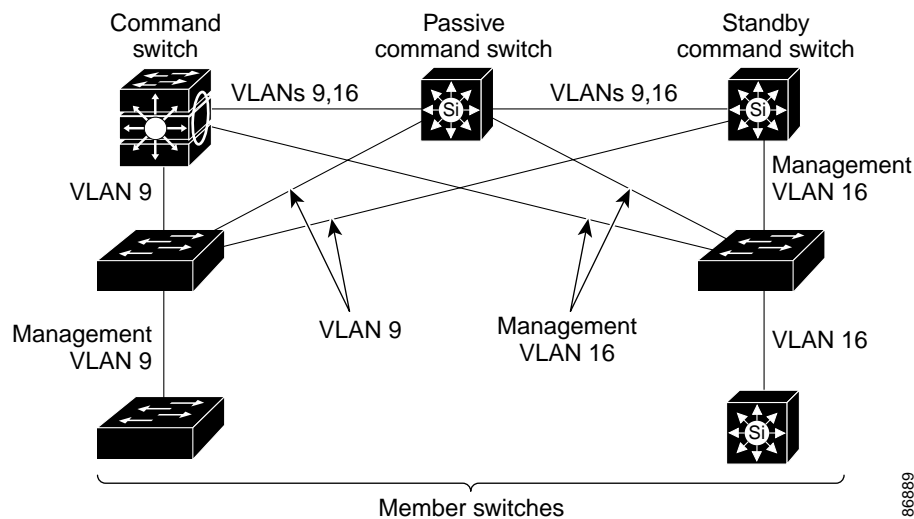
Note There is no limit to the number of switches that you can assign as standby cluster command switches. However, the total number of switches in the cluster—which would include the active cluster command switch, standby-group members, and cluster member switches—cannot be more than 16.

- Each standby-group member (Figure 5-6) must be connected to the cluster command switch through the same VLAN. Each standby-group member must also be redundantly connected to each other through at least one VLAN in common with the switch cluster.

Catalyst 1900, Catalyst 2820, Catalyst 2900 XL, Catalyst 2950, and Catalyst 3500 XL cluster member switches must be connected to the cluster standby group through their management VLANs. For more information about VLANs in switch clusters, see these sections:

- “Discovery Through Different VLANs” section on page 5-7
- “Discovery Through Different Management VLANs” section on page 5-8

Figure 5-6 VLAN Connectivity between Standby-Group Members and Cluster Members



Automatic Recovery of Cluster Configuration

The active cluster command switch continually forwards cluster-configuration information (but not device-configuration information) to the standby cluster command switch. This ensures that the standby cluster command switch can take over the cluster immediately after the active cluster command switch fails.

Automatic discovery has these limitations:

- This limitation applies only to clusters that have Catalyst 2950, Catalyst 3550, and Catalyst 3750 command and standby cluster command switches: If the active cluster command switch and standby cluster command switch become disabled *at the same time*, the passive cluster command switch with the highest priority becomes the active cluster command switch. However, because it was a passive standby cluster command switch, the previous cluster command switch *did not* forward cluster-configuration information to it. The active cluster command switch only forwards cluster-configuration information to the standby cluster command switch. You must therefore rebuild the cluster.
- This limitation applies to all clusters: If the active cluster command switch fails and there are more than two switches in the cluster standby group, the new cluster command switch does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must re-add these cluster member switches to the cluster.
- This limitation applies to all clusters: If the active cluster command switch fails and becomes active again, it does not discover any Catalyst 1900, Catalyst 2820, and Catalyst 2916M XL cluster member switches. You must again add these cluster member switches to the cluster.

When the previously active cluster command switch resumes its active role, it receives a copy of the latest cluster configuration from the active cluster command switch, including members that were added while it was down. The active cluster command switch sends a copy of the cluster configuration to the cluster standby group.

IP Addresses

You must assign IP information to a cluster command switch. You can assign more than one IP address to the cluster command switch, and you can access the cluster through any of the command-switch IP addresses. If you configure a cluster standby group, you must use the standby-group virtual IP address to manage the cluster from the active cluster command switch. Using the virtual IP address ensures that you retain connectivity to the cluster if the active cluster command switch fails and that a standby cluster command switch becomes the active cluster command switch.

If the active cluster command switch fails and the standby cluster command switch takes over, you must either use the standby-group virtual IP address or any of the IP addresses available on the new active cluster command switch to access the cluster.

You can assign an IP address to a cluster-capable switch, but it is not necessary. A cluster member switch is managed and communicates with other cluster member switches through the command-switch IP address. If the cluster member switch leaves the cluster and it does not have its own IP address, you then must assign IP information to it to manage it as a standalone switch.



Note

Changing the cluster command switch IP address ends your CMS session on the switch. Restart your CMS session by entering the new IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Internet Explorer), as described in the release notes.

For more information about IP addresses, see [Chapter 4, “Assigning the Switch IP Address and Default Gateway.”](#)

Host Names

You do not need to assign a host name to either a cluster command switch or an eligible cluster member. However, a host name assigned to the cluster command switch can help to identify the switch cluster. The default host name for the switch is *Switch*.

If a switch joins a cluster and it does not have a host name, the cluster command switch appends a unique member number to its own host name and assigns it sequentially as each switch joins the cluster. The number means the order in which the switch was added to the cluster. For example, a cluster command switch named *eng-cluster* could name the fifth cluster member *eng-cluster-5*.

If a switch has a host name, it retains that name when it joins a cluster. It retains that host name even after it leaves the cluster.

If a switch received its host name from the cluster command switch, was removed from a cluster, was then added to a new cluster, and kept the same member number (such as 5), the old host name (such as *eng-cluster-5*) is overwritten with the host name of the cluster command switch in the new cluster (such as *mkg-cluster-5*). If the switch member number changes in the new cluster (such as 3), the switch retains the previous name (*eng-cluster-5*).

Passwords

You do not need to assign passwords to an individual switch if it will be a cluster member. When a switch joins a cluster, it inherits the command-switch password and retains it when it leaves the cluster. If no command-switch password is configured, the cluster member switch inherits a null password. Cluster member switches only inherit the command-switch password.

If you change the member-switch password to be different from the command-switch password and save the change, the switch is not manageable by the cluster command switch until you change the member-switch password to match the command-switch password. Rebooting the member switch does not revert the password back to the command-switch password. We recommend that you do not change the member-switch password after it joins a cluster.

For more information about passwords, see the [“Preventing Unauthorized Access to Your Switch” section on page 7-1](#).

For password considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

SNMP Community Strings

A cluster member switch inherits the command-switch first read-only (RO) and read-write (RW) community strings with *@esN* appended to the community strings:

- *command-switch-readonly-community-string@esN*, where *N* is the member-switch number.
- *command-switch-readwrite-community-string@esN*, where *N* is the member-switch number.

If the cluster command switch has multiple read-only or read-write community strings, only the first read-only and read-write strings are propagated to the cluster member switch.

The switches support an unlimited number of community strings and string lengths. For more information about SNMP and community strings, see [Chapter 25, “Configuring SNMP.”](#)

For SNMP considerations specific to the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides specific to those switches.

TACACS+ and RADIUS

Inconsistent authentication configurations in switch clusters cause CMS to continually prompt for a user name and password. If Terminal Access Controller Access Control System Plus (TACACS+) is configured on a cluster member, it must be configured on all cluster members. Similarly, if Remote Authentication Dial-In User Service (RADIUS) is configured on a cluster member, it must be configured on all cluster members. Further, the same switch cluster cannot have some members configured with TACACS+ and other members configured with RADIUS.

For more information about TACACS+, see the [“Controlling Switch Access with TACACS+” section on page 7-10](#). For more information about RADIUS, see the [“Controlling Switch Access with RADIUS” section on page 7-18](#).

Access Modes in CMS

If your cluster has these cluster member switches running earlier software releases and if you have read-only access to these cluster member switches, some configuration windows for those switches display incomplete information:

- Catalyst 2900 XL or Catalyst 3500 XL cluster member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Catalyst 2950 cluster member switches running Cisco IOS Release 12.0(5)WC2 or earlier
- Catalyst 3550 cluster member switches running Cisco IOS Release 12.1(6)EA1 or earlier

These switches do not support read-only mode on CMS:

- Catalyst 1900 and Catalyst 2820
- Catalyst 2900 XL switches with 4-MB CPU DRAM

In read-only mode, these switches appear as unavailable devices and cannot be configured from CMS. For more information about CMS access modes, see the [“Access to Older Switches In a Cluster”](#) section on page 3-7.

LRE Profiles

A configuration conflict occurs if a switch cluster has Long-Reach Ethernet (LRE) switches that use both private and public profiles. If one LRE switch in a cluster is assigned a public profile, all LRE switches in that cluster must have that same public profile. Before you add an LRE switch to a cluster, make sure that you assign it the same public profile used by other LRE switches in the cluster.

A cluster can have a mix of LRE switches that use different private profiles.

Availability of Switch-Specific Features in Switch Clusters

The menu bar on the cluster command switch displays all options available from the switch cluster. Therefore, features specific to a cluster member switch are available from the command-switch menu bar. For example, **Device > LRE Profile** appears in the command-switch menu bar when at least one Catalyst 2900 LRE XL switch is in the cluster.

Creating a Switch Cluster

Using CMS to create a cluster is easier than using the CLI commands. This section provides this information:

- [Enabling a Cluster Command Switch, page 5-16](#)
- [Adding Cluster Member Switches, page 5-16](#)
- [Creating a Cluster Standby Group, page 5-18](#)

This section assumes you have already connected the switches, as described in the switch hardware installation guide, and followed the guidelines described in the [“Planning a Switch Cluster”](#) section on page 5-4.

**Note**

Refer to the release notes for the list of Catalyst switches eligible for switch clustering, including which ones can be cluster command switches and which ones can only be cluster member switches, and for the required software versions and browser and Java plug-in configurations.

Enabling a Cluster Command Switch

The switch you designate as the cluster command switch must meet the requirements described in the “[Cluster Command Switch Characteristics](#)” section on page 5-3, the “[Planning a Switch Cluster](#)” section on page 5-4, and the release notes.

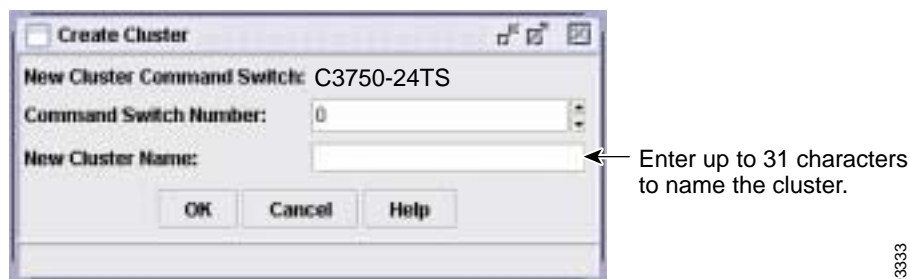
**Note**

If your switch cluster has a Catalyst 2970 switch, it should be the cluster command switch unless the cluster has a Catalyst 3750 switch or switch stack. If the switch cluster has a Catalyst 3750 switch or switch stack, that switch or switch stack must be the cluster command switch.

You can enable a cluster command switch, name the cluster, and assign an IP address and a password to the cluster command switch when you run the setup program during initial switch setup. For information about using the setup program, refer to the release notes.

If you did not enable a cluster command switch during initial switch setup, launch Device Manager from a command-capable switch, and select **Cluster > Create Cluster**. Enter a cluster number (the default is 0), and use up to 31 characters to name the cluster ([Figure 5-7](#)). Instead of using CMS to enable a cluster command switch, you can use the **cluster enable** global configuration command.

Figure 5-7 Create Cluster Window



93333

Adding Cluster Member Switches

As explained in the “[Automatic Discovery of Cluster Candidates and Members](#)” section on page 5-5, the cluster command switch automatically discovers candidate switches. When you add new cluster-capable switches to the network, the cluster command switch discovers them and adds them to a list of candidate switches.

To display an updated cluster candidates list from the Add to Cluster window ([Figure 5-8](#)), either relaunch CMS and redisplay this window, or follow these steps:

1. Close the Add to Cluster window.
2. Select **View > Refresh**.
3. Select **Cluster > Add to Cluster** to redisplay the Add to Cluster window.

From CMS, there are two ways to add switches to a cluster:

- Select **Cluster > Add to Cluster**, select a candidate switch from the list, click **Add**, and click **OK**. To add more than one candidate switch, press **Ctrl**, and make your choices, or press **Shift**, and choose the first and last switch in a range.
- Display the Topology view, right-click a candidate-switch icon, and select **Add to Cluster** (Figure 5-9). In the Topology view, candidate switches are cyan, and cluster member switches are green. To add more than one candidate switch, press **Ctrl**, and left-click the candidates that you want to add.

Instead of using CMS to add members to the cluster, you can use the **cluster member** global configuration command from the cluster command switch. Use the **password** option in this command if the candidate switch has a password.

You can select 1 or more switches as long as the total number of switches in the cluster does not exceed 16 (this includes the cluster command switch). When a cluster has 16 members, the **Add to Cluster** option is not available for that cluster. In this case, you must remove a cluster member switch before adding a new one.

If a password has been configured on a candidate switch, you are prompted to enter it before it can be added to the cluster. If the candidate switch does not have a password, any entry is ignored.

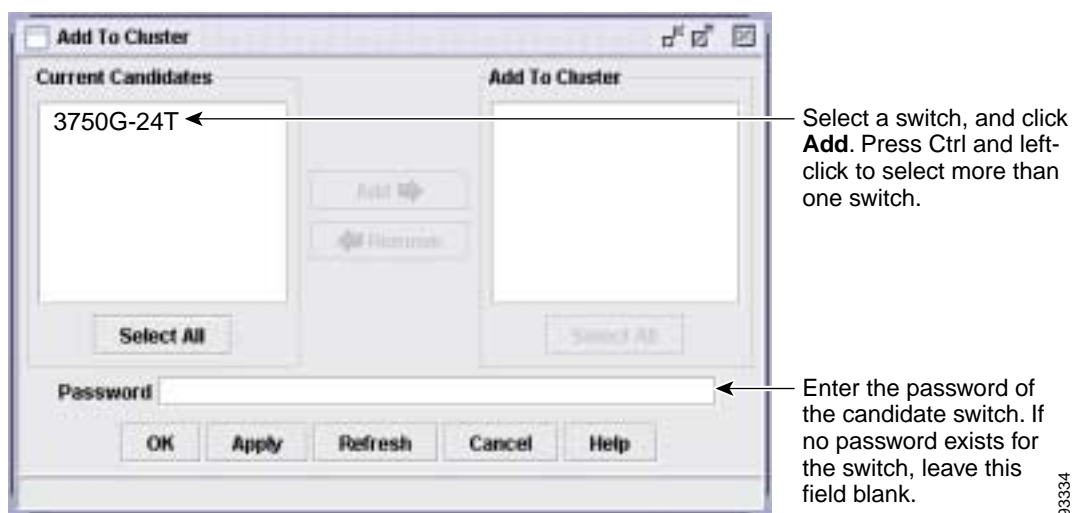
If multiple candidate switches have the same password, you can select them as a group, and add them at the same time.

If a candidate switch in the group has a password different from the group, only that specific candidate switch is not added to the cluster.

When a candidate switch joins a cluster, it inherits the command-switch password. For more information about setting passwords, see the “[Passwords](#)” section on page 5-14.

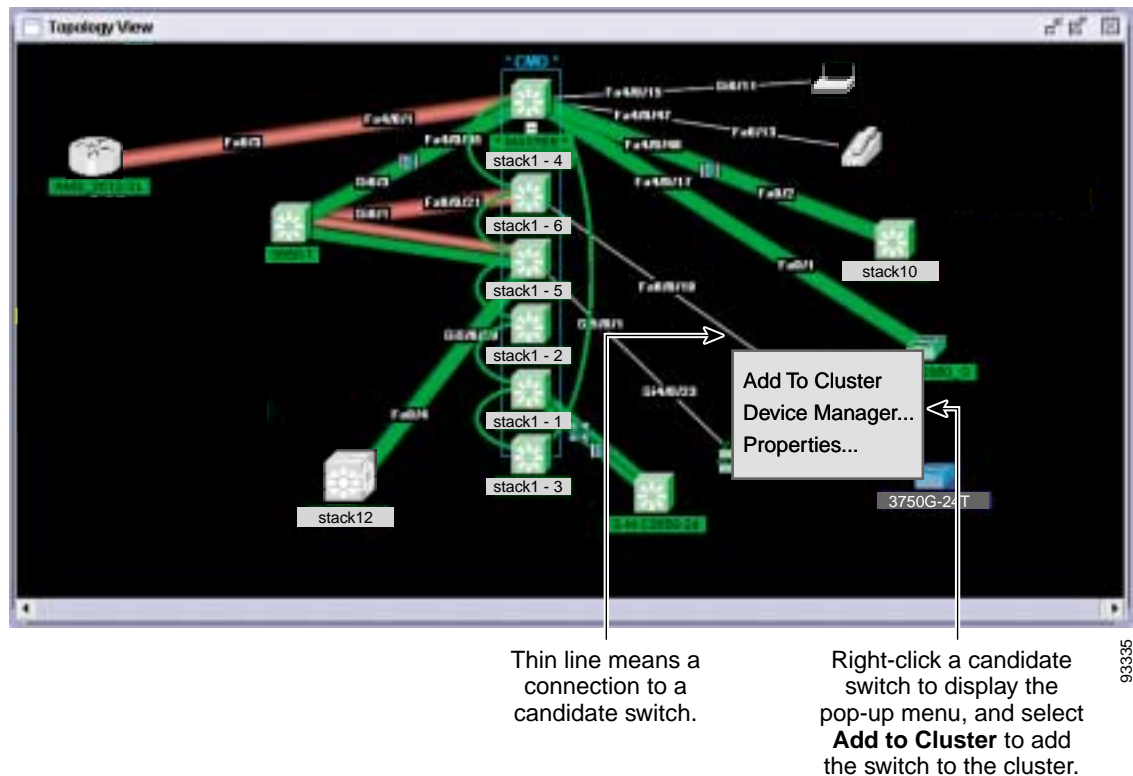
For additional authentication considerations in switch clusters, see the “[TACACS+ and RADIUS](#)” section on page 5-14.

Figure 5-8 Add to Cluster Window



93334

Figure 5-9 Using the Topology View to Add Cluster Member Switches



Creating a Cluster Standby Group

The cluster standby group members must meet the requirements described in the [“Standby Cluster Command Switch Characteristics”](#) section on page 5-3 and [“HSRP and Standby Cluster Command Switches”](#) section on page 5-10. To create a cluster standby group, select **Cluster > Standby Command Switches** (Figure 5-10).

Instead of using CMS to add switches to a standby group and to bind the standby group to a cluster, you can use the **standby ip**, the **standby name**, and the **standby priority** interface configuration commands and the **cluster standby group** global configuration command.



Note

Standby cluster command switches must be the same type of switches as the cluster command switch. For example, if the cluster command switch is a Catalyst 2970 switch, the standby cluster command switches must also be Catalyst 2970 switches. Refer to the switch configuration guide of other cluster-capable switches for their requirements on standby cluster command switches.

These abbreviations are appended to the switch host names in the Standby Command Group list to show their eligibility or status in the cluster standby group:

- AC—Active cluster command switch
- SC—Standby cluster command switch
- PC—Member of the cluster standby group but not the standby cluster command switch
- HC—Candidate switch that can be added to the cluster standby group

- CC—Cluster command switch

You must enter a virtual IP address for the cluster standby group. This address must be in the same subnet as the IP addresses of the switch. The group number must be unique within the IP subnet. It can be from 0 to 255, and the default is 0. The group name can have up to 31 characters.

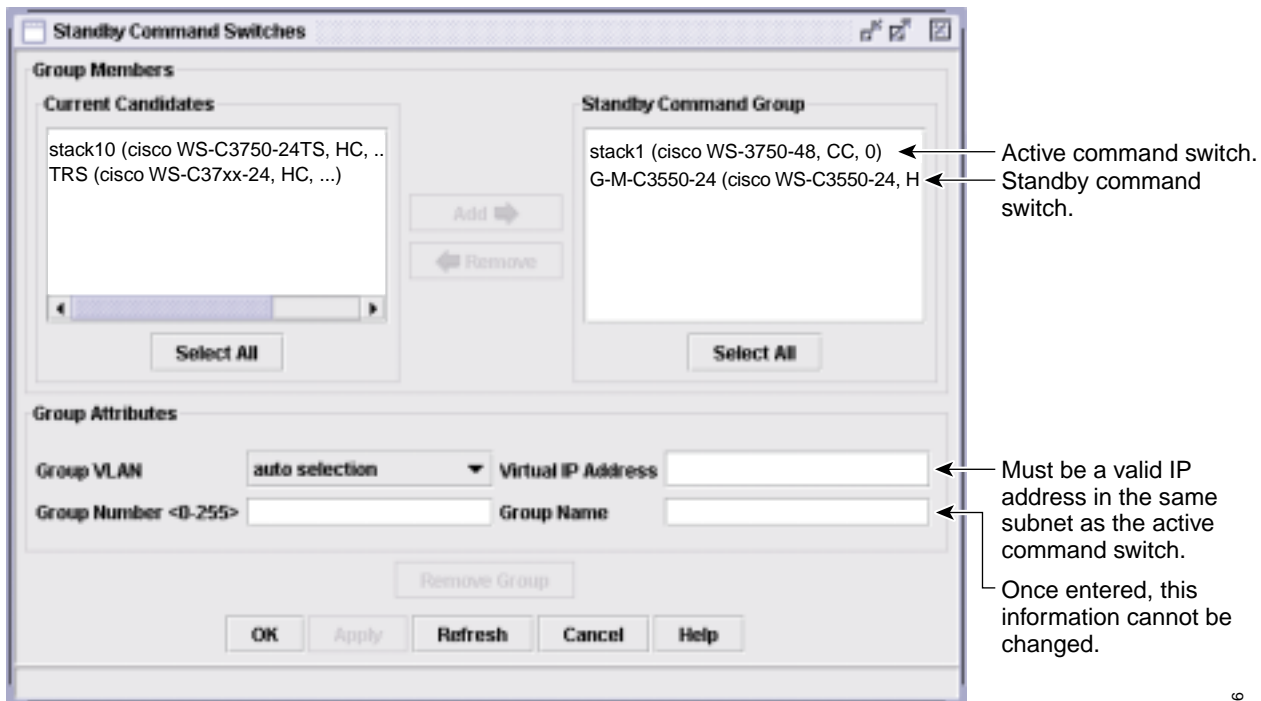
The Standby Command Configuration window uses the default values for the **preempt** and **name** commands that you have set by using the CLI. If you use this window to create the standby group, all switches in the group have the **preempt** command enabled. You must also provide a name for the group.



Note

The HSRP standby hold time interval should be greater than or equal to three times the hello time interval. The default HSRP standby hold time interval is 10 seconds. The default HSRP standby hello time interval is 3 seconds. For more information about the standby hold time and standby hello time intervals, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

Figure 5-10 Standby Command Configuration Window



93336

Verifying a Switch Cluster

When you finish adding cluster members, follow these steps to verify the cluster:

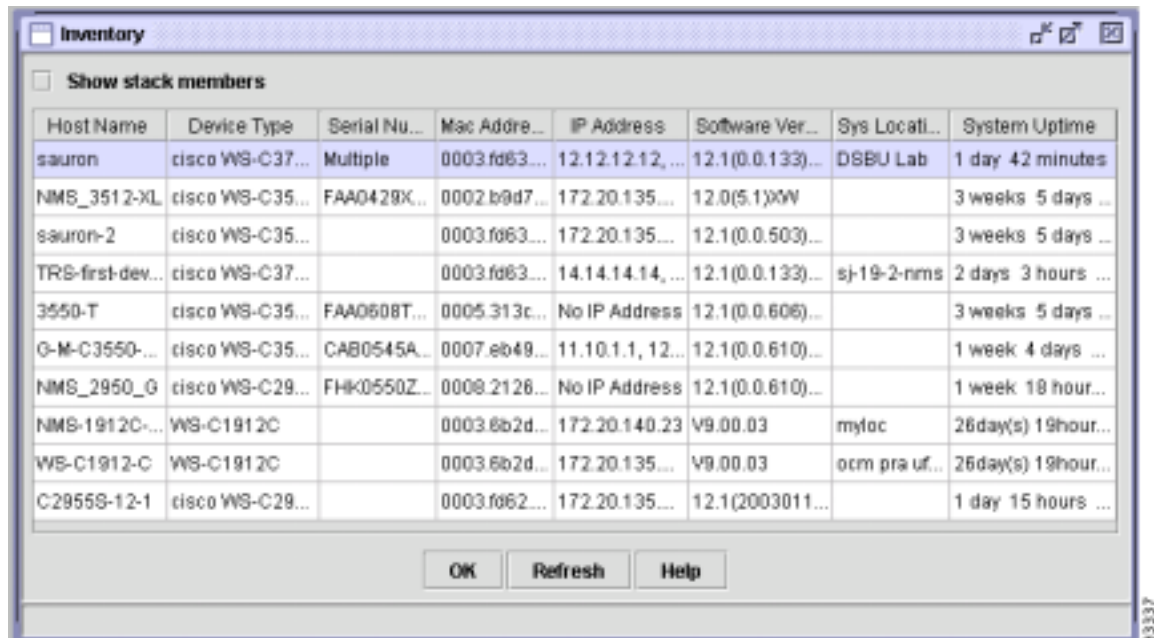
- Step 1** Enter the cluster command switch IP address in the browser **Location** field (Netscape Communicator) or **Address** field (Microsoft Internet Explorer) to access all switches in the cluster.
- Step 2** Enter the command-switch password.
- Step 3** Select **View > Topology** to display the cluster topology and to view link information (Figure 3-8 on page 3-14). For complete information about the Topology view, including descriptions of the icons, links, and colors, see the “Topology View” section on page 3-2.
- Step 4** Select **Reports > Inventory** to display an inventory of the switches in the cluster (Figure 5-11).

The summary includes information such as switch model numbers, serial numbers, software versions, IP information, and location.

You can also display port and switch statistics from **Reports > Port Statistics** and **Port > Port Settings > Runtime Status**.

Instead of using CMS to verify the cluster, you can use the **show cluster members** user EXEC command from the cluster command switch or use the **show cluster** user EXEC command from the cluster command switch or from a cluster member switch.

Figure 5-11 Inventory Window



The screenshot shows a window titled "Inventory" with a checkbox for "Show stack members". Below it is a table with the following columns: Host Name, Device Type, Serial Nu..., Mac Addre..., IP Address, Software Ver..., Sys Locati..., and System Uptime. The table contains 11 rows of switch data.

Host Name	Device Type	Serial Nu...	Mac Addre...	IP Address	Software Ver...	Sys Locati...	System Uptime
sauron	cisco WS-C37...	Multiple	0003.f063...	12.12.12.12...	12.1(0.0.133)...	DSBU Lab	1 day 42 minutes
NMS_3512-XL	cisco WS-C35...	FAA0429X...	0002.b9d7...	172.20.135...	12.0(5.1)XW		3 weeks 5 days ...
sauron-2	cisco WS-C35...		0003.f063...	172.20.135...	12.1(0.0.503)...		3 weeks 5 days ...
TRB-first-dev...	cisco WS-C37...		0003.f063...	14.14.14.14...	12.1(0.0.133)...	sj-19-2-nms	2 days 3 hours ...
3550-T	cisco WS-C35...	FAA0508T...	0005.313c...	No IP Address	12.1(0.0.606)...		3 weeks 5 days ...
G-M-C3550...	cisco WS-C35...	CAB0545A...	0007.eb49...	11.10.1.1, 12...	12.1(0.0.610)...		1 week 4 days ...
NMS_2950_G	cisco WS-C29...	FHK0550Z...	0008.2126...	No IP Address	12.1(0.0.610)...		1 week 18 hour...
NMS-1912C...	WS-C1912C		0003.6b2d...	172.20.140.23	V9.00.03	myloc	26day(s) 19hour...
WS-C1912-C	WS-C1912C		0003.6b2d...	172.20.135...	V9.00.03	ocm pra uf...	26day(s) 19hour...
C2955G-12-1	cisco WS-C29...		0003.f062...	172.20.135...	12.1(2003011...		1 day 15 hours ...

At the bottom of the window are buttons for OK, Refresh, and Help.

If you lose connectivity with a cluster member switch or if a cluster command switch fails, see the cluster-related recovery procedures in Chapter 29, “Troubleshooting.”

For more information about creating and managing clusters, refer to the online help. For information about the cluster commands, refer to the switch command reference.

Using the CLI to Manage Switch Clusters

You can configure cluster member switches from the CLI by first logging into the cluster command switch. Enter the **rcommand** user EXEC command and the cluster member switch number to start a Telnet session (through a console or Telnet connection) and to access the cluster member switch CLI. The command mode changes, and the Cisco IOS commands operate as usual. Enter the **exit** privileged EXEC command on the cluster member switch to return to the command-switch CLI.

This example shows how to log into member-switch 3 from the command-switch CLI:

```
switch# rcommand 3
```

If you do not know the member-switch number, enter the **show cluster members** privileged EXEC command on the cluster command switch. For more information about the **rcommand** command and all other cluster commands, refer to the switch command reference.

The Telnet session accesses the member-switch CLI at the same privilege level as on the cluster command switch. The Cisco IOS commands then operate as usual. For instructions on configuring the switch for a Telnet session, see the [“Disabling Password Recovery” section on page 7-5](#).

Catalyst 1900 and Catalyst 2820 CLI Considerations

If your switch cluster has Catalyst 1900 and Catalyst 2820 switches running standard edition software, the Telnet session accesses the management console (a menu-driven interface) if the cluster command switch is at privilege level 15. If the cluster command switch is at privilege level 1 to 14, you are prompted for the password to access the menu console.

Command-switch privilege levels map to the Catalyst 1900 and Catalyst 2820 cluster member switches running standard and Enterprise Edition Software as follows:

- If the command-switch privilege level is 1 to 14, the cluster member switch is accessed at privilege level 1.
- If the command-switch privilege level is 15, the cluster member switch is accessed at privilege level 15.



Note

The Catalyst 1900 and Catalyst 2820 CLI is available only on switches running Enterprise Edition Software.

For more information about the Catalyst 1900 and Catalyst 2820 switches, refer to the installation and configuration guides for those switches.

Using SNMP to Manage Switch Clusters

When you first power on the switch, SNMP is enabled if you enter the IP information by using the setup program and accept its proposed configuration. If you did not use the setup program to enter the IP information and SNMP was not enabled, you can enable it as described in the [“Configuring SNMP” section on page 25-6](#). On Catalyst 1900 and Catalyst 2820 switches, SNMP is enabled by default.

When you create a cluster, the cluster command switch manages the exchange of messages between cluster member switches and an SNMP application. The cluster software on the cluster command switch appends the cluster member switch number (*@esN*, where *N* is the switch number) to the first configured read-write and read-only community strings on the cluster command switch and propagates them to the

cluster member switch. The cluster command switch uses this community string to control the forwarding of gets, sets, and get-next messages between the SNMP management station and the cluster member switches.

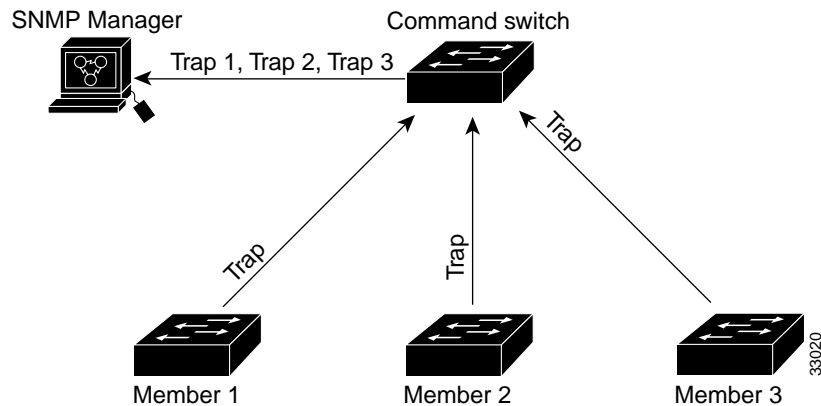
**Note**

When a cluster standby group is configured, the cluster command switch can change without your knowledge. Use the first read-write and read-only community strings to communicate with the cluster command switch if there is a cluster standby group configured for the cluster.

If the cluster member switch does not have an IP address, the cluster command switch redirects traps from the cluster member switch to the management station, as shown in [Figure 5-12](#). If a cluster member switch has its own IP address and community strings, the cluster member switch can send traps directly to the management station, without going through the cluster command switch.

If a cluster member switch has its own IP address and community strings, they can be used in addition to the access provided by the cluster command switch. For more information about SNMP and community strings, see [Chapter 25, "Configuring SNMP."](#)

Figure 5-12 *SNMP Management for a Cluster*





Administering the Switch

This chapter describes how to perform one-time operations to administer the Catalyst 2970 switch.

This chapter consists of these sections:

- [Managing the System Time and Date, page 6-1](#)
- [Configuring a System Name and Prompt, page 6-16](#)
- [Creating a Banner, page 6-19](#)
- [Managing the MAC Address Table, page 6-22](#)
- [Managing the ARP Table, page 6-29](#)

Managing the System Time and Date

You can manage the system time and date on your switch using automatic configuration, such as the Network Time Protocol (NTP), or manual configuration methods.



Note

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- [Understanding the System Clock, page 6-2](#)
- [Understanding Network Time Protocol, page 6-2](#)
- [Configuring NTP, page 6-4](#)
- [Configuring Time and Date Manually, page 6-11](#)

Understanding the System Clock

The heart of the time service is the system clock. This clock runs from the moment the system starts up and keeps track of the date and time.

The system clock can then be set from these sources:

- Network Time Protocol
- Manual configuration

The system clock can provide time to these services:

- User **show** commands
- Logging and debugging messages

The system clock keeps track of time internally based on Universal Time Coordinated (UTC), also known as Greenwich Mean Time (GMT). You can configure information about the local time zone and summer time (daylight saving time) so that the time correctly appears for the local time zone.

The system clock keeps track of whether the time is *authoritative* or not (that is, whether it has been set by a time source considered to be authoritative). If it is not authoritative, the time is available only for display purposes and is not redistributed. For configuration information, see the [“Configuring Time and Date Manually” section on page 6-11](#).

Understanding Network Time Protocol

The NTP is designed to time-synchronize a network of devices. NTP runs over User Datagram Protocol (UDP), which runs over IP. NTP is documented in RFC 1305.

An NTP network usually gets its time from an authoritative time source, such as a radio clock or an atomic clock attached to a time server. NTP then distributes this time across the network. NTP is extremely efficient; no more than one packet per minute is necessary to synchronize two devices to within a millisecond of one another.

NTP uses the concept of a *stratum* to describe how many NTP hops away a device is from an authoritative time source. A stratum 1 time server has a radio or atomic clock directly attached, a stratum 2 time server receives its time through NTP from a stratum 1 time server, and so on. A device running NTP automatically chooses as its time source the device with the lowest stratum number with which it communicates through NTP. This strategy effectively builds a self-organizing tree of NTP speakers.

NTP avoids synchronizing to a device whose time might not be accurate by never synchronizing to a device that is not synchronized. NTP also compares the time reported by several devices and does not synchronize to a device whose time is significantly different than the others, even if its stratum is lower.

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP address of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, in that case, information flow is one-way only.

The time kept on a device is a critical resource; you should use the security features of NTP to avoid the accidental or malicious setting of an incorrect time. Two mechanisms are available: an access list-based restriction scheme and an encrypted authentication mechanism.

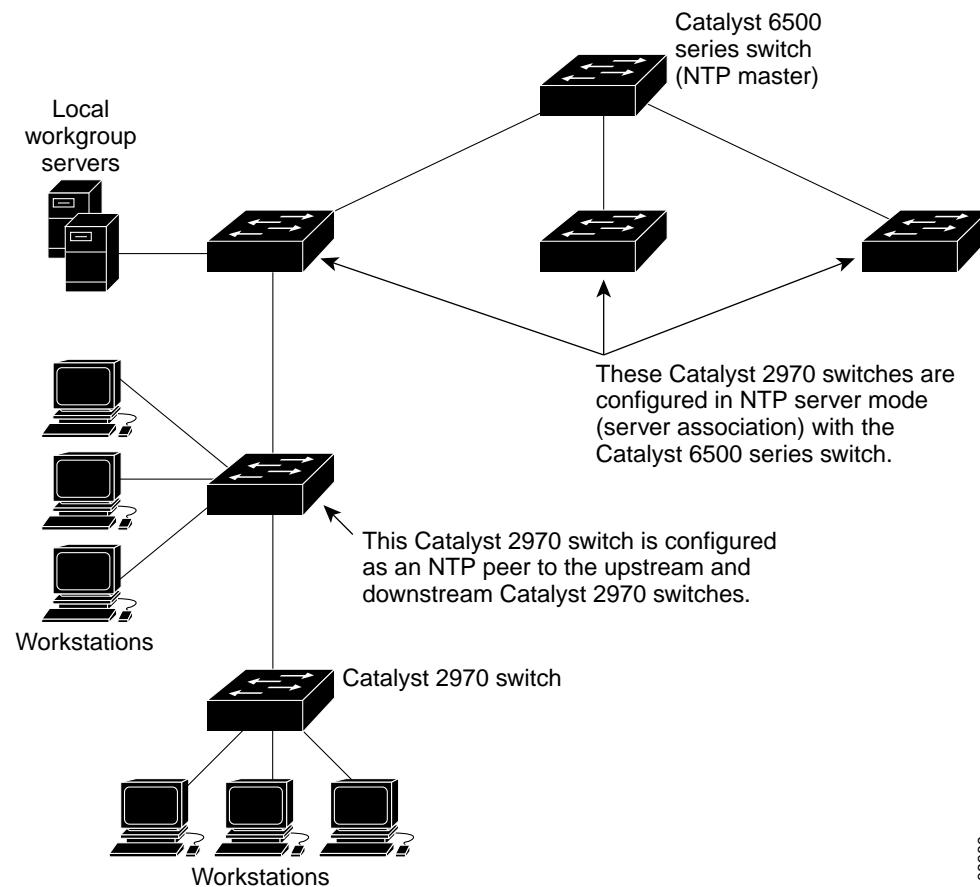
Cisco's implementation of NTP does not support stratum 1 service; it is not possible to connect to a radio or atomic clock. We recommend that the time service for your network be derived from the public NTP servers available on the IP Internet. [Figure 6-1](#) show a typical network example using NTP.

If the network is isolated from the Internet, Cisco's implementation of NTP allows a device to act as though it is synchronized through NTP, when in fact it has determined the time by using other means. Other devices then synchronize to that device through NTP.

When multiple sources of time are available, NTP is always considered to be more authoritative. NTP time overrides the time set by any other method.

Several manufacturers include NTP software for their host systems, and a publicly available version for systems running UNIX and its various derivatives is also available. This software allows host systems to be time-synchronized as well.

Figure 6-1 Typical NTP Network Configuration



89389

Configuring NTP

The switch does not have a hardware-supported clock and cannot function as an NTP master clock to which peers synchronize themselves when an external NTP source is not available. The switch also has no hardware support for a calendar. As a result, the **ntp update-calendar** and the **ntp master** global configuration commands are not available.

This section contains this configuration information:

- [Default NTP Configuration, page 6-4](#)
- [Configuring NTP Authentication, page 6-5](#)
- [Configuring NTP Associations, page 6-6](#)
- [Configuring NTP Broadcast Service, page 6-7](#)
- [Configuring NTP Access Restrictions, page 6-8](#)
- [Configuring the Source IP Address for NTP Packets, page 6-10](#)
- [Displaying the NTP Configuration, page 6-11](#)

Default NTP Configuration

[Table 6-1](#) shows the default NTP configuration.

Table 6-1 *Default NTP Configuration*

Feature	Default Setting
NTP authentication	Disabled. No authentication key is specified.
NTP peer or server associations	None configured.
NTP broadcast service	Disabled; no interface sends or receives NTP broadcast packets.
NTP access restrictions	No access control is specified.
NTP packet source IP address	The source address is determined by the outgoing interface.

NTP is enabled on all interfaces by default. All interfaces receive NTP packets.

Configuring NTP Authentication

This procedure must be coordinated with the administrator of the NTP server; the information you configure in this procedure must be matched by the servers used by the switch to synchronize its time to the NTP server.

Beginning in privileged EXEC mode, follow these steps to authenticate the associations (communications between devices running NTP that provide for accurate timekeeping) with other devices for security purposes:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp authenticate	Enable the NTP authentication feature, which is disabled by default.
Step 3	ntp authentication-key <i>number</i> md5 <i>value</i>	<p>Define the authentication keys. By default, none are defined.</p> <ul style="list-style-type: none"> For <i>number</i>, specify a key number. The range is 1 to 4294967295. md5 specifies that message authentication support is provided by using the message digest algorithm 5 (MD5). For <i>value</i>, enter an arbitrary string of up to eight characters for the key. <p>The switch does not synchronize to a device unless both have one of these authentication keys, and the key number is specified by the ntp trusted-key <i>key-number</i> command.</p>
Step 4	ntp trusted-key <i>key-number</i>	<p>Specify one or more key numbers (defined in Step 3) that a peer NTP device must provide in its NTP packets for this switch to synchronize to it.</p> <p>By default, no trusted keys are defined.</p> <p>For <i>key-number</i>, specify the key defined in Step 3.</p> <p>This command provides protection against accidentally synchronizing the switch to a device that is not trusted.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable NTP authentication, use the **no ntp authenticate** global configuration command. To remove an authentication key, use the **no ntp authentication-key** *number* global configuration command. To disable authentication of the identity of a device, use the **no ntp trusted-key** *key-number* global configuration command.

This example shows how to configure the switch to synchronize only to devices providing authentication key 42 in the device's NTP packets:

```
Switch(config)# ntp authenticate
Switch(config)# ntp authentication-key 42 md5 aNiceKey
Switch(config)# ntp trusted-key 42
```

Configuring NTP Associations

An NTP association can be a peer association (this switch can either synchronize to the other device or allow the other device to synchronize to it), or it can be a server association (meaning that only this switch synchronizes to the other device, and not the other way around).

Beginning in privileged EXEC mode, follow these steps to form an NTP association with another device:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp peer <i>ip-address</i> [version <i>number</i>] [key <i>keyid</i>] [source <i>interface</i>] [prefer] or ntp server <i>ip-address</i> [version <i>number</i>] [key <i>keyid</i>] [source <i>interface</i>] [prefer]	Configure the switch system clock to synchronize a peer or to be synchronized by a peer (peer association). or Configure the switch system clock to be synchronized by a time server (server association). No peer or server associations are defined by default. <ul style="list-style-type: none"> For <i>ip-address</i> in a peer association, specify either the IP address of the peer providing, or being provided, the clock synchronization. For a server association, specify the IP address of the time server providing the clock synchronization. (Optional) For <i>number</i>, specify the NTP version number. The range is 1 to 3. By default, version 3 is selected. (Optional) For <i>keyid</i>, enter the authentication key defined with the ntp authentication-key global configuration command. (Optional) For <i>interface</i>, specify the interface from which to pick the IP source address. By default, the source IP address is taken from the outgoing interface. (Optional) Enter the prefer keyword to make this peer or server the preferred one that provides synchronization. This keyword reduces switching back and forth between peers and servers.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

You need to configure only one end of an association; the other device can automatically establish the association. If you are using the default NTP version (version 3) and NTP synchronization does not occur, try using NTP version 2. Many NTP servers on the Internet run version 2.

To remove a peer or server association, use the **no ntp peer** *ip-address* or the **no ntp server** *ip-address* global configuration command.

This example shows how to configure the switch to synchronize its system clock with the clock of the peer at IP address 172.16.22.44 using NTP version 2:

```
Switch(config)# ntp server 172.16.22.44 version 2
```

Configuring NTP Broadcast Service

The communications between devices running NTP (known as *associations*) are usually statically configured; each device is given the IP addresses of all devices with which it should form associations. Accurate timekeeping is possible by exchanging NTP messages between each pair of devices with an association. However, in a LAN environment, NTP can be configured to use IP broadcast messages instead. This alternative reduces configuration complexity because each device can simply be configured to send or receive broadcast messages. However, the information flow is one-way only.

The switch can send or receive NTP broadcast packets on an interface-by-interface basis if there is an NTP broadcast server, such as a router, broadcasting time information on the network. The switch can send NTP broadcast packets to a peer so that the peer can synchronize to it. The switch can also receive NTP broadcast packets to synchronize its own clock. This section provides procedures for both sending and receiving NTP broadcast packets.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send NTP broadcast packets to peers so that they can synchronize their clock to the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to send NTP broadcast packets, and enter interface configuration mode.
Step 3	ntp broadcast [version <i>number</i>] [key <i>keyid</i>] [<i>destination-address</i>]	Enable the interface to send NTP broadcast packets to a peer. By default, this feature is disabled on all interfaces. <ul style="list-style-type: none"> • (Optional) For <i>number</i>, specify the NTP version number. The range is 1 to 3. If you do not specify a version, version 3 is used. • (Optional) For <i>keyid</i>, specify the authentication key to use when sending packets to the peer. • (Optional) For <i>destination-address</i>, specify the IP address of the peer that is synchronizing its clock to this switch.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 7		Configure the connected peers to receive NTP broadcast packets as described in the next procedure.

To disable the interface from sending NTP broadcast packets, use the **no ntp broadcast** interface configuration command.

This example shows how to configure an interface to send NTP version 2 packets:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ntp broadcast version 2
```

Beginning in privileged EXEC mode, follow these steps to configure the switch to receive NTP broadcast packets from connected peers:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to receive NTP broadcast packets, and enter interface configuration mode.
Step 3	ntp broadcast client	Enable the interface to receive NTP broadcast packets. By default, no interfaces receive NTP broadcast packets.
Step 4	exit	Return to global configuration mode.
Step 5	ntp broadcastdelay <i>microseconds</i>	(Optional) Change the estimated round-trip delay between the switch and the NTP broadcast server. The default is 3000 microseconds; the range is 1 to 999999.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an interface from receiving NTP broadcast packets, use the **no ntp broadcast client** interface configuration command. To change the estimated round-trip delay to the default, use the **no ntp broadcastdelay** global configuration command.

This example shows how to configure an interface to receive NTP broadcast packets:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ntp broadcast client
```

Configuring NTP Access Restrictions

You can control NTP access on two levels as described in these sections:

- [Creating an Access Group and Assigning a Basic IP Access List, page 6-9](#)
- [Disabling NTP Services on a Specific Interface, page 6-10](#)

Creating an Access Group and Assigning a Basic IP Access List

Beginning in privileged EXEC mode, follow these steps to control access to NTP services by using access lists:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp access-group { query-only serve-only serve peer } access-list-number	<p>Create an access group, and apply a basic IP access list.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • query-only—Allows only NTP control queries. • serve-only—Allows only time requests. • serve—Allows time requests and NTP control queries, but does not allow the switch to synchronize to the remote device. • peer—Allows time requests and NTP control queries and allows the switch to synchronize to the remote device. <p>For <i>access-list-number</i>, enter a standard IP access list number from 1 to 99.</p>
Step 3	access-list access-list-number permit source [source-wildcard]	<p>Create the access list.</p> <ul style="list-style-type: none"> • For <i>access-list-number</i>, enter the number specified in Step 2. • Enter the permit keyword to permit access if the conditions are matched. • For <i>source</i>, enter the IP address of the device that is permitted access to the switch. • (Optional) For <i>source-wildcard</i>, enter the wildcard bits to be applied to the source. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The access group keywords are scanned in this order, from least restrictive to most restrictive:

1. **peer**—Allows time requests and NTP control queries and allows the switch to synchronize itself to a device whose address passes the access list criteria.
2. **serve**—Allows time requests and NTP control queries, but does not allow the switch to synchronize itself to a device whose address passes the access list criteria.
3. **serve-only**—Allows only time requests from a device whose address passes the access list criteria.
4. **query-only**—Allows only NTP control queries from a device whose address passes the access list criteria.

If the source IP address matches the access lists for more than one access type, the first type is granted. If no access groups are specified, all access types are granted to all devices. If any access groups are specified, only the specified access types are granted.

To remove access control to the switch NTP services, use the **no ntp access-group {query-only | serve-only | serve | peer}** global configuration command.

This example shows how to configure the switch to allow itself to synchronize to a peer from access list 99. However, the switch restricts access to allow only time requests from access list 42:

```
Switch# configure terminal
Switch(config)# ntp access-group peer 99
Switch(config)# ntp access-group serve-only 42
Switch(config)# access-list 99 permit 172.20.130.5
Switch(config)# access list 42 permit 172.20.130.6
```

Disabling NTP Services on a Specific Interface

NTP services are enabled on all interfaces by default.

Beginning in privileged EXEC mode, follow these steps to disable NTP packets from being received on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Enter interface configuration mode, and specify the interface to disable.
Step 3	ntp disable	Disable NTP packets from being received on the interface. By default, all interfaces receive NTP packets.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable receipt of NTP packets on an interface, use the **no ntp disable** interface configuration command.

Configuring the Source IP Address for NTP Packets

When the switch sends an NTP packet, the source IP address is normally set to the address of the interface through which the NTP packet is sent. Use the **ntp source** global configuration command when you want to use a particular source IP address for all NTP packets. The address is taken from the specified interface. This command is useful if the address on an interface cannot be used as the destination for reply packets.

Beginning in privileged EXEC mode, follow these steps to configure a specific interface from which the IP source address is to be taken:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ntp source <i>type number</i>	Specify the interface type and number from which the IP source address is taken. By default, the source address is determined by the outgoing interface.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The specified interface is used for the source address for all packets sent to all destinations. If a source address is to be used for a specific association, use the **source** keyword in the **ntp peer** or **ntp server** global configuration command as described in the [“Configuring NTP Associations”](#) section on page 6-6.

Displaying the NTP Configuration

You can use two privileged EXEC commands to display NTP information:

- **show ntp associations [detail]**
- **show ntp status**

For detailed information about the fields in these displays, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

Configuring Time and Date Manually

If no other source of time is available, you can manually configure the time and date after the system is restarted. The time remains accurate until the next system restart. We recommend that you use manual configuration only as a last resort. If you have an outside source to which the switch can synchronize, you do not need to manually set the system clock.

This section contains this configuration information:

- [Setting the System Clock, page 6-12](#)
- [Displaying the Time and Date Configuration, page 6-12](#)
- [Configuring the Time Zone, page 6-13](#)
- [Configuring Summer Time \(Daylight Saving Time\), page 6-14](#)

Setting the System Clock

If you have an outside source on the network that provides time services, such as an NTP server, you do not need to manually set the system clock.

Beginning in privileged EXEC mode, follow these steps to set the system clock:

	Command	Purpose
Step 1	clock set <i>hh:mm:ss day month year</i> or clock set <i>hh:mm:ss month day year</i>	Manually set the system clock using one of these formats. <ul style="list-style-type: none">For <i>hh:mm:ss</i>, specify the time in hours (24-hour format), minutes, and seconds. The time specified is relative to the configured time zone.For <i>day</i>, specify the day by date in the month.For <i>month</i>, specify the month by name.For <i>year</i>, specify the year (no abbreviation).
Step 2	show running-config	Verify your entries.
Step 3	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to manually set the system clock to 1:32 p.m. on July 23, 2001:

```
Switch# clock set 13:32:00 23 July 2001
```

Displaying the Time and Date Configuration

To display the time and date configuration, use the **show clock [detail]** privileged EXEC command.

The system clock keeps an *authoritative* flag that shows whether the time is authoritative (believed to be accurate). If the system clock has been set by a timing source such as NTP, the flag is set. If the time is not authoritative, it is used only for display purposes. Until the clock is authoritative and the *authoritative* flag is set, the flag prevents peers from synchronizing to the clock when the peers' time is invalid.

The symbol that precedes the **show clock** display has this meaning:

- *—Time is not authoritative.
- (blank)—Time is authoritative.
- .—Time is authoritative, but NTP is not synchronized.

Configuring the Time Zone

Beginning in privileged EXEC mode, follow these steps to manually configure the time zone:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clock timezone <i>zone hours-offset</i> [<i>minutes-offset</i>]	Set the time zone. The switch keeps internal time in universal time coordinated (UTC), so this command is used only for display purposes and when the time is manually set. <ul style="list-style-type: none"> For <i>zone</i>, enter the name of the time zone to be displayed when standard time is in effect. The default is UTC. For <i>hours-offset</i>, enter the hours offset from UTC. (Optional) For <i>minutes-offset</i>, enter the minutes offset from UTC.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The *minutes-offset* variable in the **clock timezone** global configuration command is available for those cases where a local time zone is a percentage of an hour different from UTC. For example, the time zone for some sections of Atlantic Canada (AST) is UTC-3.5, where the 3 means 3 hours and .5 means 50 percent. In this case, the necessary command is **clock timezone AST -3 30**.

To set the time to UTC, use the **no clock timezone** global configuration command.

Configuring Summer Time (Daylight Saving Time)

Beginning in privileged EXEC mode, follow these steps to configure summer time (daylight saving time) in areas where it starts and ends on a particular day of the week each year:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clock summer-time zone recurring [<i>week day month hh:mm week day month hh:mm [offset]</i>]	<p>Configure summer time to start and end on the specified days every year. Summer time is disabled by default. If you specify clock summer-time zone recurring without parameters, the summer time rules default to the United States rules.</p> <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

This example shows how to specify that summer time starts on the first Sunday in April at 02:00 and ends on the last Sunday in October at 02:00:

```
Switch(config)# clock summer-time PDT recurring 1 Sunday April 2:00 last Sunday October 2:00
```

Beginning in privileged EXEC mode, follow these steps if summer time in your area does not follow a recurring pattern (configure the exact date and time of the next summer time events):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clock summer-time zone date [<i>month date year hh:mm month date year hh:mm [offset]</i>] or clock summer-time zone date [<i>date month year hh:mm date month year hh:mm [offset]</i>]	Configure summer time to start on the first date and end on the second date. Summer time is disabled by default. <ul style="list-style-type: none"> For <i>zone</i>, specify the name of the time zone (for example, PDT) to be displayed when summer time is in effect. (Optional) For <i>week</i>, specify the week of the month (1 to 5 or last). (Optional) For <i>day</i>, specify the day of the week (Sunday, Monday...). (Optional) For <i>month</i>, specify the month (January, February...). (Optional) For <i>hh:mm</i>, specify the time (24-hour format) in hours and minutes. (Optional) For <i>offset</i>, specify the number of minutes to add during summer time. The default is 60.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The first part of the **clock summer-time** global configuration command specifies when summer time begins, and the second part specifies when it ends. All times are relative to the local time zone. The start time is relative to standard time. The end time is relative to summer time. If the starting month is after the ending month, the system assumes that you are in the southern hemisphere.

To disable summer time, use the **no clock summer-time** global configuration command.

This example shows how to set summer time to start on October 12, 2000, at 02:00, and end on April 26, 2001, at 02:00:

```
Switch(config)# clock summer-time pdt date 12 October 2000 2:00 26 April 2001 2:00
```

Configuring a System Name and Prompt

You configure the system name on the switch to identify it. By default, the system name and prompt are *Switch*.

If you have not configured a system prompt, the first 20 characters of the system name are used as the system prompt. A greater-than symbol [>] is appended. The prompt is updated whenever the system name changes, unless you manually configure the prompt by using the **prompt** global configuration command.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference* and the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This section contains this configuration information:

- [Default System Name and Prompt Configuration, page 6-16](#)
- [Configuring a System Name, page 6-16](#)
- [Configuring a System Prompt, page 6-17](#)
- [Understanding DNS, page 6-17](#)

Default System Name and Prompt Configuration

The default switch system name and prompt is *Switch*.

Configuring a System Name

Beginning in privileged EXEC mode, follow these steps to manually configure a system name:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	hostname <i>name</i>	Manually configure a system name. The default setting is <i>switch</i> . The name must follow the rules for ARPANET host names. They must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and hyphens. Names can be up to 63 characters.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you set the system name, it is also used as the system prompt. You can override the prompt setting by using the **prompt** global configuration command.

To return to the default hostname, use the **no hostname** global configuration command.

Configuring a System Prompt

Beginning in privileged EXEC mode, follow these steps to manually configure a system prompt:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	prompt <i>string</i>	Configure the command-line prompt to override the setting from the hostname command. The default prompt is either <i>switch</i> or the name defined with the hostname global configuration command, followed by an angle bracket (>) for user EXEC mode or a pound sign (#) for privileged EXEC mode. The prompt can consist of all printing characters and escape sequences.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default prompt, use the **no prompt** [*string*] global configuration command.

Understanding DNS

The DNS protocol controls the Domain Name System (DNS), a distributed database with which you can map host names to IP addresses. When you configure DNS on your switch, you can substitute the host name for the IP address with all IP commands, such as **ping**, **telnet**, **connect**, and related Telnet support operations.

IP defines a hierarchical naming scheme that allows a device to be identified by its location or domain. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that IP identifies by a *com* domain name, so its domain name is *cisco.com*. A specific device in this domain, for example, the File Transfer Protocol (FTP) system is identified as *ftp.cisco.com*.

To keep track of domain names, IP has defined the concept of a domain name server, which holds a cache (or database) of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, specify the name server that is present on your network, and enable the DNS.

This section contains this configuration information:

- [Default DNS Configuration, page 6-18](#)
- [Setting Up DNS, page 6-18](#)
- [Displaying the DNS Configuration, page 6-19](#)

Default DNS Configuration

Table 6-2 shows the default DNS configuration.

Table 6-2 Default DNS Configuration

Feature	Default Setting
DNS enable state	Enabled.
DNS default domain name	None configured.
DNS servers	No name server addresses are configured.

Setting Up DNS

Beginning in privileged EXEC mode, follow these steps to set up your switch to use the DNS:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip domain-name <i>name</i>	<p>Define a default domain name that the software uses to complete unqualified host names (names without a dotted-decimal domain name).</p> <p>Do not include the initial period that separates an unqualified name from the domain name.</p> <p>At boot time, no domain name is configured; however, if the switch configuration comes from a BOOTP or Dynamic Host Configuration Protocol (DHCP) server, then the default domain name might be set by the BOOTP or DHCP server (if the servers were configured with this information).</p>
Step 3	ip name-server <i>server-address1</i> [<i>server-address2</i> ... <i>server-address6</i>]	<p>Specify the address of one or more name servers to use for name and address resolution.</p> <p>You can specify up to six name servers. Separate each server address with a space. The first server specified is the primary server. The switch sends DNS queries to the primary server first. If that query fails, the backup servers are queried.</p>
Step 4	ip domain-lookup	<p>(Optional) Enable DNS-based host name-to-address translation on your switch. This feature is enabled by default.</p> <p>If your network devices require connectivity with devices in networks for which you do not control name assignment, you can dynamically assign device names that uniquely identify your devices by using the global Internet naming scheme (DNS).</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If you use the switch IP address as its hostname, the IP address is used and no DNS query occurs. If you configure a hostname that contains no periods (.), a period followed by the default domain name is appended to the hostname before the DNS query is made to map the name to an IP address. The default domain name is the value set by the **ip domain-name** global configuration command. If there is a period (.) in the hostname, the Cisco IOS software looks up the IP address without appending any default domain name to the hostname.

To remove a domain name, use the **no ip domain-name** *name* global configuration command. To remove a name server address, use the **no ip name-server** *server-address* global configuration command. To disable DNS on the switch, use the **no ip domain-lookup** global configuration command.

Displaying the DNS Configuration

To display the DNS configuration information, use the **show running-config** privileged EXEC command.

Creating a Banner

You can configure a message-of-the-day (MOTD) and a login banner. The MOTD banner displays on all connected terminals at login and is useful for sending messages that affect all network users (such as impending system shutdowns).

The login banner also displays on all connected terminals. It appears after the MOTD banner and before the login prompts.



Note

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This section contains this configuration information:

- [Default Banner Configuration, page 6-19](#)
- [Configuring a Message-of-the-Day Login Banner, page 6-20](#)
- [Configuring a Login Banner, page 6-21](#)

Default Banner Configuration

The MOTD and login banners are not configured.

Configuring a Message-of-the-Day Login Banner

You can create a single or multiline message banner that appears on the screen when someone logs in to the switch.

Beginning in privileged EXEC mode, follow these steps to configure a MOTD login banner:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	banner motd <i>c message c</i>	Specify the message of the day. For <i>c</i> , enter the delimiting character of your choice, for example, a pound sign (#), and press the Return key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For <i>message</i> , enter a banner message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the MOTD banner, use the **no banner motd** global configuration command.

This example shows how to configure a MOTD banner for the switch by using the pound sign (#) symbol as the beginning and ending delimiter:

```
Switch(config)# banner motd #
This is a secure site. Only authorized users are allowed.
For access, contact technical support.
#
Switch(config)#
```

This example shows the banner that appears from the previous configuration:

```
Unix> telnet 172.2.5.4
Trying 172.2.5.4...
Connected to 172.2.5.4.
Escape character is '^]'.
```

```
This is a secure site. Only authorized users are allowed.
For access, contact technical support.
```

```
User Access Verification
```

```
Password:
```


Configuring a Login Banner

You can configure a login banner to be displayed on all connected terminals. This banner appears after the MOTD banner and before the login prompt.

Beginning in privileged EXEC mode, follow these steps to configure a login banner:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	banner login <i>c message c</i>	Specify the login message. For <i>c</i> , enter the delimiting character of your choice, for example, a pound sign (#), and press the Return key. The delimiting character signifies the beginning and end of the banner text. Characters after the ending delimiter are discarded. For <i>message</i> , enter a login message up to 255 characters. You cannot use the delimiting character in the message.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the login banner, use the **no banner login** global configuration command.

This example shows how to configure a login banner for the switch by using the dollar sign (\$) symbol as the beginning and ending delimiter:

```
Switch(config)# banner login $
Access for authorized users only. Please enter your username and password.
$
Switch(config)#
```

Managing the MAC Address Table

The MAC address table contains address information that the switch uses to forward traffic between ports. All MAC addresses in the address table are associated with one or more ports. The address table includes these types of addresses:

- **Dynamic address:** a source MAC address that the switch learns and then ages when it is not in use.
- **Static address:** a manually entered unicast address that does not age and that is not lost when the switch resets.

The address table lists the destination MAC address, the associated VLAN ID, and port number associated with the address and the type (static or dynamic).

**Note**

For complete syntax and usage information for the commands used in this section, refer to the command reference for this release.

This section contains this configuration information:

- [Building the Address Table, page 6-22](#)
- [MAC Addresses and VLANs, page 6-23](#)
- [Default MAC Address Table Configuration, page 6-23](#)
- [Changing the Address Aging Time, page 6-23](#)
- [Removing Dynamic Address Entries, page 6-24](#)
- [Configuring MAC Address Notification Traps, page 6-24](#)
- [Adding and Removing Static Address Entries, page 6-26](#)
- [Configuring Unicast MAC Address Filtering, page 6-27](#)
- [Displaying Address Table Entries, page 6-28](#)

Building the Address Table

With multiple MAC addresses supported on all ports, you can connect any port on the switch to individual workstations, repeaters, switches, routers, or other network devices. The switch provides dynamic addressing by learning the source address of packets it receives on each port and adding the address and its associated port number to the address table. As stations are added or removed from the network, the switch updates the address table, adding new dynamic addresses and aging out those that are not in use.

The aging interval is globally configured. However, the switch maintains an address table for each VLAN, and STP can accelerate the aging interval on a per-VLAN basis.

The switch sends packets between any combination of ports, based on the destination address of the received packet. Using the MAC address table, the switch forwards the packet only to the port associated with the destination address. If the destination address is on the port that sent the packet, the packet is filtered and not forwarded. The switch always uses the store-and-forward method: complete packets are stored and checked for errors before transmission.

MAC Addresses and VLANs

All addresses are associated with a VLAN. An address can exist in more than one VLAN and have different destinations in each. Unicast addresses, for example, could forward to port 1 in VLAN 1 and port 9 in VLAN 5.



Note

Multiport static addresses are not supported.

Each VLAN maintains its own logical address table. A known address in one VLAN is unknown in another until it is learned or statically associated with a port in the other VLAN.

Default MAC Address Table Configuration

Table 6-3 shows the default MAC address table configuration.

Table 6-3 Default MAC Address Table Configuration

Feature	Default Setting
Aging time	300 seconds
Dynamic addresses	Automatically learned
Static addresses	None configured

Changing the Address Aging Time

Dynamic addresses are source MAC addresses that the switch learns and then ages when they are not in use. You can change the aging time setting for all VLANs or for a specified VLAN.

Setting too short an aging time can cause addresses to be prematurely removed from the table. Then when the switch receives a packet for an unknown destination, it floods the packet to all ports in the same VLAN as the receiving port. This unnecessary flooding can impact performance. Setting too long an aging time can cause the address table to be filled with unused addresses, which prevents new addresses from being learned. Flooding results, which can impact switch performance.

Beginning in privileged EXEC mode, follow these steps to configure the dynamic address table aging time:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table aging-time [0 10-1000000] [vlan <i>vlan-id</i>]	Set the length of time that a dynamic entry remains in the MAC address table after the entry is used or updated. The range is 10 to 1000000 seconds. The default is 300. You can also enter 0, which disables aging. Static address entries are never aged or removed from the table. For <i>vlan-id</i> , valid IDs are 1 to 4094. Do not enter leading zeros.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table aging-time	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no mac address-table aging-time** global configuration command.

Removing Dynamic Address Entries

To remove all dynamic entries, use the **clear mac address-table dynamic** command in privileged EXEC mode. You can also remove a specific MAC address (**clear mac address-table dynamic address** *mac-address*), remove all addresses on the specified physical port or port channel (**clear mac address-table dynamic interface** *interface-id*), or remove all addresses on a specified VLAN (**clear mac address-table dynamic vlan** *vlan-id*).

To verify that dynamic entries have been removed, use the **show mac address-table dynamic** privileged EXEC command.

Configuring MAC Address Notification Traps

MAC address notification enables you to track users on a network by storing the MAC address activity on the switch. Whenever the switch learns or removes a MAC address, an SNMP notification can be generated and sent to the NMS. If you have many users coming and going from the network, you can set a trap interval time to bundle the notification traps and reduce network traffic. The MAC notification history table stores the MAC address activity for each hardware port for which the trap is enabled. MAC address notifications are generated for dynamic and secure MAC addresses; events are not generated for self addresses, multicast addresses, or other static addresses.

Beginning in privileged EXEC mode, follow these steps to configure the switch to send MAC address notification traps to an NMS host:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server host <i>host-addr</i> { traps informs } { version { 1 2c 3 }} <i>community-string</i> <i>notification-type</i>	Specify the recipient of the trap message. <ul style="list-style-type: none"> For <i>host-addr</i>, specify the name or address of the NMS. Specify traps (the default) to send SNMP traps to the host. Specify informs to send SNMP informs to the host. Specify the SNMP version to support. Version 1, the default, is not available with informs. For <i>community-string</i>, specify the string to send with the notification operation. Though you can set this string by using the snmp-server host command, we recommend that you define this string by using the snmp-server community command before using the snmp-server host command. For <i>notification-type</i>, use the mac-notification keyword.
Step 3	snmp-server enable traps mac-notification	Enable the switch to send MAC address traps to the NMS.
Step 4	mac address-table notification	Enable the MAC address notification feature.
Step 5	mac address-table notification [interval <i>value</i>] [history-size <i>value</i>]	Enter the trap interval time and the history table size. <ul style="list-style-type: none"> (Optional) For interval <i>value</i>, specify the notification trap interval in seconds between each set of traps that are generated to the NMS. The range is 0 to 2147483647 seconds; the default is 1 second. (Optional) For history-size <i>value</i>, specify the maximum number of entries in the MAC notification history table. The range is 0 to 500; the default is 1.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to enable the SNMP MAC address notification trap.
Step 7	snmp trap mac-notification { added removed }	Enable the MAC address notification trap. <ul style="list-style-type: none"> Enable the MAC notification trap whenever a MAC address is added on this interface. Enable the MAC notification trap whenever a MAC address is removed from this interface.
Step 8	end	Return to privileged EXEC mode.

	Command	Purpose
Step 9	show mac address-table notification interface show running-config	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the switch from sending MAC address notification traps, use the **no snmp-server enable traps mac-notification** global configuration command. To disable the MAC address notification traps on a specific interface, use the **no snmp trap mac-notification {added | removed}** interface configuration command. To disable the MAC address notification feature, use the **no mac address-table notification** global configuration command.

This example shows how to specify 172.20.10.10 as the NMS, enable the switch to send MAC address notification traps to the NMS, enable the MAC address notification feature, set the interval time to 60 seconds, set the history-size to 100 entries, and enable traps whenever a MAC address is added on Gigabit Ethernet interface 0/4.

```
Switch(config)# snmp-server host 172.20.10.10 traps private
Switch(config)# snmp-server enable traps mac-notification
Switch(config)# mac address-table notification
Switch(config)# mac address-table notification interval 60
Switch(config)# mac address-table notification history-size 100
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# snmp trap mac-notification added
```

You can verify the previous commands by entering the **show mac address-table notification interface** and the **show mac address-table notification** privileged EXEC commands.

Adding and Removing Static Address Entries

A static address has these characteristics:

- It is manually entered in the address table and must be manually removed.
- It can be a unicast address.
- It does not age and is retained when the switch restarts.

You can add and remove static addresses and define the forwarding behavior for them. The forwarding behavior determines how a port that receives a packet forwards it to another port for transmission. Because all ports are associated with at least one VLAN, the switch acquires the VLAN ID for the address from the ports that you specify.

A packet with a static address that arrives on a VLAN where it has not been statically entered is flooded to all ports and not learned.

You add a static address to the address table by specifying the destination MAC unicast address and the VLAN from which it is received. Packets received with this destination address are forwarded to the interface specified with the *interface-id* option.

Beginning in privileged EXEC mode, follow these steps to add a static address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table static <i>mac-addr</i> vlan <i>vlan-id</i> interface <i>interface-id</i>	Add a static address to the MAC address table. <ul style="list-style-type: none"> For <i>mac-addr</i>, specify the destination MAC unicast address to add to the address table. Packets with this destination address received in the specified VLAN are forwarded to the specified interface. For <i>vlan-id</i>, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 4094; do not enter leading zeros. For <i>interface-id</i>..., specify the interface to which the received packet is forwarded. Valid interfaces include physical ports.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table static	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove static entries from the address table, use the **no mac address-table static** *mac-addr* **vlan** *vlan-id* [**interface** *interface-id*] global configuration command.

This example shows how to add the static address c2f3.220a.12f4 to the MAC address table. When a packet is received in VLAN 4 with this MAC address as its destination address, the packet is forwarded to the specified interface:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 interface
gigabitethernet0/1
```

Configuring Unicast MAC Address Filtering

When unicast MAC address filtering is enabled, the switch drops packets with specific source or destination MAC addresses. This feature is disabled by default and only supports unicast static addresses.

Follow these guidelines when using this feature:

- Multicast MAC addresses, broadcast MAC addresses, and router MAC addresses are not supported. If you specify one of these addresses when entering the **mac address-table static** *mac-addr* **vlan** *vlan-id* **drop** global configuration command, one of these messages appears:

```
% Only unicast addresses can be configured to be dropped
% CPU destined address cannot be configured as drop address
```
- Packets that are forwarded to the CPU are also not supported.
- If you add a unicast MAC address as a static address and configure unicast MAC address filtering, the switch either adds the MAC address as a static address or drops packets with that MAC address, depending on which command was entered last. The second command that you entered overrides the first command.

For example, if you enter the **mac address-table static *mac-addr* vlan *vlan-id* interface *interface-id*** global configuration command followed by the **mac address-table static *mac-addr* vlan *vlan-id* drop** command, the switch drops packets with the specified MAC address as a source or destination.

If you enter the **mac address-table static *mac-addr* vlan *vlan-id* drop** global configuration command followed by the **mac address-table static *mac-addr* vlan *vlan-id* interface *interface-id*** command, the switch adds the MAC address as a static address.

You enable unicast MAC address filtering and configure the switch to drop packets with a specific address by specifying the source or destination unicast MAC address and the VLAN from which it is received.

Beginning in privileged EXEC mode, follow these steps to configure the switch to drop a source or destination unicast static address:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac address-table static <i>mac-addr</i> vlan <i>vlan-id</i> drop	Enable unicast MAC address filtering and configure the switch to drop a packet with the specified source or destination unicast static address. <ul style="list-style-type: none"> For <i>mac-addr</i>, specify a source or destination unicast MAC address. Packets with this MAC address are dropped. For <i>vlan-id</i>, specify the VLAN for which the packet with the specified MAC address is received. Valid VLAN IDs are 1 to 4094.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mac address-table static	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable unicast MAC address filtering, use the **no mac address-table static *mac-addr* vlan *vlan-id*** global configuration command.

This example shows how to enable unicast MAC address filtering and to configure the switch to drop packets that have a source or destination address of c2f3.220a.12f4. When a packet is received in VLAN 4 with this MAC address as its source or destination, the packet is dropped:

```
Switch(config)# mac address-table static c2f3.220a.12f4 vlan 4 drop
```

Displaying Address Table Entries

You can display the MAC address table by using one or more of the privileged EXEC commands described in [Table 6-4](#):

Table 6-4 Commands for Displaying the MAC Address Table

Command	Description
show mac address-table address	Displays MAC address table information for the specified MAC address.
show mac address-table aging-time	Displays the aging time in all VLANs or the specified VLAN.
show mac address-table count	Displays the number of addresses present in all VLANs or the specified VLAN.
show mac address-table dynamic	Displays only dynamic MAC address table entries.

Table 6-4 Commands for Displaying the MAC Address Table (continued)

Command	Description
show mac address-table interface	Displays the MAC address table information for the specified interface.
show mac address-table multicast	Displays the Layer 2 multicast entries for all VLANs or the specified VLAN.
show mac address-table notification	Displays the MAC notification parameters and history table.
show mac address-table static	Displays only static MAC address table entries.
show mac address-table vlan	Displays the MAC address table information for the specified VLAN.

Managing the ARP Table

To communicate with a device (over Ethernet, for example), the software first must determine the 48-bit MAC or the local data link address of that device. The process of determining the local data link address from an IP address is called *address resolution*.

The Address Resolution Protocol (ARP) associates a host IP address with the corresponding media or MAC addresses and the VLAN ID. Taking an IP address as input, ARP determines the associated MAC address. Once a MAC address is determined, the IP-MAC address association is stored in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests and replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP). By default, standard Ethernet-style ARP encapsulation (represented by the **arpa** keyword) is enabled on the IP interface.

ARP entries added manually to the table do not age and must be manually removed.

For CLI procedures, refer to the Cisco IOS Release 12.1 documentation on Cisco.com.



Configuring Switch-Based Authentication

This chapter describes how to configure switch-based authentication on the Catalyst 2970 switch.

This chapter consists of these sections:

- [Preventing Unauthorized Access to Your Switch, page 7-1](#)
- [Protecting Access to Privileged EXEC Commands, page 7-2](#)
- [Controlling Switch Access with TACACS+, page 7-10](#)
- [Controlling Switch Access with RADIUS, page 7-18](#)
- [Controlling Switch Access with Kerberos, page 7-32](#)
- [Configuring the Switch for Local Authentication and Authorization, page 7-36](#)
- [Configuring the Switch for Secure Shell, page 7-37](#)

Preventing Unauthorized Access to Your Switch

You can prevent unauthorized users from reconfiguring your switch and viewing configuration information. Typically, you want network administrators to have access to your switch while you restrict access to users who dial from outside the network through an asynchronous port, connect from outside the network through a serial port, or connect through a terminal or workstation from within the local network.

To prevent unauthorized access into your switch, you should configure one or more of these security features:

- At a minimum, you should configure passwords and privileges at each switch port. These passwords are locally stored on the switch. When users attempt to access the switch through a port or line, they must enter the password specified for the port or line before they can access the switch. For more information, see the [“Protecting Access to Privileged EXEC Commands” section on page 7-2](#).
- For an additional layer of security, you can also configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair. For more information, see the [“Configuring Username and Password Pairs” section on page 7-7](#).
- If you want to use username and password pairs, but you want to store them centrally on a server instead of locally, you can store them in a database on a security server. Multiple networking devices can then use the same database to obtain user authentication (and, if necessary, authorization) information. For more information, see the [“Controlling Switch Access with TACACS+” section on page 7-10](#).

Protecting Access to Privileged EXEC Commands

A simple way of providing terminal access control in your network is to use passwords and assign privilege levels. Password protection restricts access to a network or network device. Privilege levels define what commands users can enter after they have logged into a network device.



Note

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section describes how to control access to the configuration file and privileged EXEC commands. It contains this configuration information:

- [Default Password and Privilege Level Configuration, page 7-2](#)
- [Setting or Changing a Static Enable Password, page 7-3](#)
- [Protecting Enable and Enable Secret Passwords with Encryption, page 7-4](#)
- [Disabling Password Recovery, page 7-5](#)
- [Setting a Telnet Password for a Terminal Line, page 7-6](#)
- [Configuring Username and Password Pairs, page 7-7](#)
- [Configuring Multiple Privilege Levels, page 7-8](#)

Default Password and Privilege Level Configuration

[Table 7-1](#) shows the default password and privilege level configuration.

Table 7-1 Default Password and Privilege Levels

Feature	Default Setting
Enable password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is not encrypted in the configuration file.
Enable secret password and privilege level	No password is defined. The default is level 15 (privileged EXEC level). The password is encrypted before it is written to the configuration file.
Line password	No password is defined.

Setting or Changing a Static Enable Password

The enable password controls access to the privileged EXEC mode. Beginning in privileged EXEC mode, follow these steps to set or change a static enable password:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	enable password <i>password</i>	<p>Define a new password or change an existing password for access to privileged EXEC mode.</p> <p>By default, no password is defined.</p> <p>For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. It can contain the question mark (?) character if you precede the question mark with the key combination Ctrl-v when you create the password; for example, to create the password abc?123, do this:</p> <p>Enter abc.</p> <p>Enter Ctrl-v.</p> <p>Enter ?123.</p> <p>When the system prompts you to enter the enable password, you need not precede the question mark with the Ctrl-v; you can simply enter abc?123 at the password prompt.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	<p>(Optional) Save your entries in the configuration file.</p> <p>The enable password is not encrypted and can be read in the switch configuration file.</p>

To remove the password, use the **no enable password** global configuration command.

This example shows how to change the enable password to *11u2c3k4y5*. The password is not encrypted and provides access to level 15 (traditional privileged EXEC mode access):

```
Switch(config)# enable password 11u2c3k4y5
```

Protecting Enable and Enable Secret Passwords with Encryption

To provide an additional layer of security, particularly for passwords that cross the network or that are stored on a Trivial File Transfer Protocol (TFTP) server, you can use either the **enable password** or **enable secret** global configuration commands. Both commands accomplish the same thing; that is, you can establish an encrypted password that users must enter to access privileged EXEC mode (the default) or any privilege level you specify.

We recommend that you use the **enable secret** command because it uses an improved encryption algorithm.

If you configure the **enable secret** command, it takes precedence over the **enable password** command; the two commands cannot be in effect simultaneously.

Beginning in privileged EXEC mode, follow these steps to configure encryption for enable and enable secret passwords:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	enable password [level <i>level</i>] { <i>password</i> <i>encryption-type encrypted-password</i> } or enable secret [level <i>level</i>] { <i>password</i> <i>encryption-type encrypted-password</i> }	Define a new password or change an existing password for access to privileged EXEC mode. or Define a secret password, which is saved using a nonreversible encryption method. <ul style="list-style-type: none">• (Optional) For <i>level</i>, the range is from 0 to 15. Level 1 is normal user EXEC mode privileges. The default level is 15 (privileged EXEC mode privileges).• For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.• (Optional) For <i>encryption-type</i>, only type 5, a Cisco proprietary encryption algorithm, is available. If you specify an encryption type, you must provide an encrypted password—an encrypted password that you copy from another switch configuration. Note If you specify an encryption type and then enter a clear text password, you can not re-enter privileged EXEC mode. You cannot recover a lost encrypted password by any method.
Step 3	service password-encryption	(Optional) Encrypt the password when the password is defined or when the configuration is written. Encryption prevents the password from being readable in the configuration file.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

If both the enable and enable secret passwords are defined, users must enter the enable secret password.

Use the **level** keyword to define a password for a specific privilege level. After you specify the level and set a password, give the password only to users who need to have access at this level. Use the **privilege level** global configuration command to specify commands accessible at various levels. For more information, see the [“Configuring Multiple Privilege Levels” section on page 7-8](#).

If you enable password encryption, it applies to all passwords including username passwords, authentication key passwords, the privileged command password, and console and virtual terminal line passwords.

To remove a password and level, use the **no enable password [level level]** or **no enable secret [level level]** global configuration command. To disable password encryption, use the **no service password-encryption** global configuration command.

This example shows how to configure the encrypted password `1FaD0$Xyti5Rkls3LoyxzS8` for privilege level 2:

```
Switch(config)# enable secret level 2 5 $1$FaD0$Xyti5Rkls3LoyxzS8
```

Disabling Password Recovery

By default, any end user with physical access to the switch can recover from a lost password by interrupting the boot process while the switch is powering on and then by entering a new password.

The password-recovery disable feature protects access to the switch password by disabling part of this functionality. When this feature is enabled, the end user can interrupt the boot process only by agreeing to set the system back to the default configuration. With password recovery disabled, you can still interrupt the boot process and change the password, but the configuration file (config.text) and the VLAN database file (vlan.dat) are deleted.



Note

If you disable password recovery, we recommend that you keep a backup copy of the configuration file on a secure server in case the end user interrupts the boot process and sets the system back to default values. Do not keep a backup copy of the configuration file on the switch. If the switch is operating in VTP transparent mode, we recommend that you also keep a backup copy of the VLAN database file on a secure server. When the switch is returned to the default system configuration, you can download the saved files to the switch by using the XMODEM protocol. For more information, see the [“Recovering from a Lost or Forgotten Password” section on page 29-4](#).

Beginning in privileged EXEC mode, follow these steps to disable password recovery:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no service password-recovery	Disable password recovery. This setting is saved in an area of the Flash memory that is accessible by the boot loader and the Cisco IOS image, but it is not part of the file system and is not accessible by any user.
Step 3	end	Return to privileged EXEC mode.
Step 4	show version	Verify the configuration by checking the last few lines of the command output.

To re-enable password recovery, use the **service password-recovery** global configuration command.



Note

Disabling password recovery will not work if you have set the switch to boot manually by using the **boot manual** global configuration command. This command produces the boot loader prompt (*switch:*) after the switch is power cycled.

Setting a Telnet Password for a Terminal Line

When you power-up your switch for the first time, an automatic setup program runs to assign IP information and to create a default configuration for continued use. The setup program also prompts you to configure your switch for Telnet access through a password. If you did not configure this password during the setup program, you can configure it now through the command-line interface (CLI).

Beginning in privileged EXEC mode, follow these steps to configure your switch for Telnet access:

	Command	Purpose
Step 1		Attach a PC or workstation with emulation software to the switch console port. The default data characteristics of the console port are 9600, 8, 1, no parity. You might need to press the Return key several times to see the command-line prompt.
Step 2	enable password <i>password</i>	Enter privileged EXEC mode.
Step 3	configure terminal	Enter global configuration mode.
Step 4	line vty 0 15	Configure the number of Telnet sessions (lines), and enter line configuration mode. There are 16 possible sessions on a command-capable switch. The 0 and 15 mean that you are configuring all 16 possible Telnet sessions.
Step 5	password <i>password</i>	Enter a Telnet password for the line or lines. For <i>password</i> , specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries. The password is listed under the command line vty 0 15 .
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the password, use the **no password** global configuration command.

This example shows how to set the Telnet password to *let45me67in89*:

```
Switch(config)# line vty 10
Switch(config-line)# password let45me67in89
```


Configuring Username and Password Pairs

You can configure username and password pairs, which are locally stored on the switch. These pairs are assigned to lines or interfaces and authenticate each user before that user can access the switch. If you have defined privilege levels, you can also assign a specific privilege level (with associated rights and privileges) to each username and password pair.

Beginning in privileged EXEC mode, follow these steps to establish a username-based authentication system that requests a login username and a password:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type password</i> }	Enter the username, privilege level, and password for each user. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 1 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password will follow. Enter 7 to specify that a hidden password will follow. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
Step 3	line console 0 or line vty 0 15	Enter line configuration mode, and configure the console port (line 0) or the VTY lines (line 0 to 15).
Step 4	login local	Enable local password checking at login time. Authentication is based on the username specified in Step 2.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable username authentication for a specific user, use the **no username** *name* global configuration command. To disable password checking and allow connections without a password, use the **no login** line configuration command.

Configuring Multiple Privilege Levels

By default, the Cisco IOS software has two modes of password security: user EXEC and privileged EXEC. You can configure up to 16 hierarchical levels of commands for each mode. By configuring multiple passwords, you can allow different sets of users to have access to specified commands.

For example, if you want many users to have access to the **clear line** command, you can assign it level 2 security and distribute the level 2 password fairly widely. But if you want more restricted access to the **configure** command, you can assign it level 3 security and distribute that password to a more restricted group of users.

This section includes this configuration information:

- [Setting the Privilege Level for a Command, page 7-8](#)
- [Changing the Default Privilege Level for Lines, page 7-9](#)
- [Logging into and Exiting a Privilege Level, page 7-10](#)

Setting the Privilege Level for a Command

Beginning in privileged EXEC mode, follow these steps to set the privilege level for a command mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	privilege mode level level command	Set the privilege level for a command. <ul style="list-style-type: none"> • For <i>mode</i>, enter configure for global configuration mode, exec for EXEC mode, interface for interface configuration mode, or line for line configuration mode. • For <i>level</i>, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the enable password. • For <i>command</i>, specify the command to which you want to restrict access.
Step 3	enable password level level password	Specify the enable password for the privilege level. <ul style="list-style-type: none"> • For <i>level</i>, the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. • For <i>password</i>, specify a string from 1 to 25 alphanumeric characters. The string cannot start with a number, is case sensitive, and allows spaces but ignores leading spaces. By default, no password is defined.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config or show privilege	Verify your entries. The first command shows the password and access level configuration. The second command shows the privilege level configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When you set a command to a privilege level, all commands whose syntax is a subset of that command are also set to that level. For example, if you set the **show ip traffic** command to level 15, the **show** commands and **show ip** commands are automatically set to privilege level 15 unless you set them individually to different levels.

To return to the default privilege for a given command, use the **no privilege mode level level command** global configuration command.

This example shows how to set the **configure** command to privilege level 14 and define *SecretPswd14* as the password users must enter to use level 14 commands:

```
Switch(config)# privilege exec level 14 configure
Switch(config)# enable password level 14 SecretPswd14
```

Changing the Default Privilege Level for Lines

Beginning in privileged EXEC mode, follow these steps to change the default privilege level for a line:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line vty line	Select the virtual terminal line on which to restrict access.
Step 3	privilege level level	Change the default privilege level for the line. For <i>level</i> , the range is from 0 to 15. Level 1 is for normal user EXEC mode privileges. Level 15 is the level of access permitted by the enable password.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config or show privilege	Verify your entries. The first command shows the password and access level configuration. The second command shows the privilege level configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Users can override the privilege level you set using the **privilege level** line configuration command by logging in to the line and enabling a different privilege level. They can lower the privilege level by using the **disable** command. If users know the password to a higher privilege level, they can use that password to enable the higher privilege level. You might specify a high level or privilege level for your console line to restrict line usage.

To return to the default line privilege level, use the **no privilege level** line configuration command.

Logging into and Exiting a Privilege Level

Beginning in privileged EXEC mode, follow these steps to log in to a specified privilege level and to exit to a specified privilege level:

	Command	Purpose
Step 1	enable <i>level</i>	Log in to a specified privilege level. For <i>level</i> , the range is 0 to 15.
Step 2	disable <i>level</i>	Exit to a specified privilege level. For <i>level</i> , the range is 0 to 15.

Controlling Switch Access with TACACS+

This section describes how to enable and configure Terminal Access Controller Access Control System Plus (TACACS+), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. TACACS+ is facilitated through authentication, authorization, accounting (AAA) and can be enabled only through AAA commands.



Note

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section contains this configuration information:

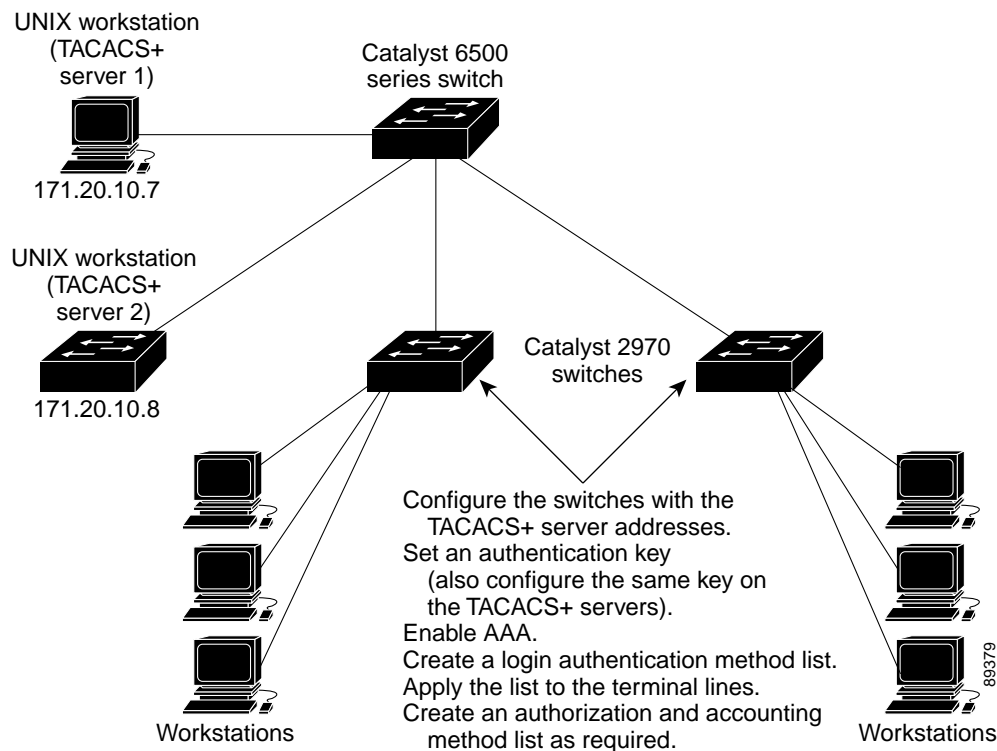
- [Understanding TACACS+, page 7-10](#)
- [TACACS+ Operation, page 7-12](#)
- [Configuring TACACS+, page 7-13](#)
- [Displaying the TACACS+ Configuration, page 7-17](#)

Understanding TACACS+

TACACS+ is a security application that provides centralized validation of users attempting to gain access to your switch. TACACS+ services are maintained in a database on a TACACS+ daemon typically running on a UNIX or Windows NT workstation. You should have access to and should configure a TACACS+ server before the configuring TACACS+ features on your switch.

TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service—authentication, authorization, and accounting—independently. Each service can be tied into its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

The goal of TACACS+ is to provide a method for managing multiple network access points from a single management service. Your switch can be a network access server along with other Cisco routers and access servers. A network access server provides connections to a single user, to a network or subnetwork, and to interconnected networks as shown in [Figure 7-1](#).

Figure 7-1 Typical TACACS+ Network Configuration

TACACS+, administered through the AAA security services, can provide these services:

- **Authentication**—Provides complete control of authentication through login and password dialog, challenge and response, and messaging support.

The authentication facility can conduct a dialog with the user (for example, after a username and password are provided, to challenge a user with several questions, such as home address, mother's maiden name, service type, and social security number). The TACACS+ authentication service can also send messages to user screens. For example, a message could notify users that their passwords must be changed because of the company's password aging policy.

- **Authorization**—Provides fine-grained control over user capabilities for the duration of the user's session, including but not limited to setting autocommands, access control, session duration, or protocol support. You can also enforce restrictions on what commands a user can execute with the TACACS+ authorization feature.
- **Accounting**—Collects and sends information used for billing, auditing, and reporting to the TACACS+ daemon. Network managers can use the accounting facility to track user activity for a security audit or to provide information for user billing. Accounting records include user identities, start and stop times, executed commands (such as PPP), number of packets, and number of bytes.

The TACACS+ protocol provides authentication between the switch and the TACACS+ daemon, and it ensures confidentiality because all protocol exchanges between the switch and the TACACS+ daemon are encrypted.

You need a system running the TACACS+ daemon software to use TACACS+ on your switch.

TACACS+ Operation

When a user attempts a simple ASCII login by authenticating to a switch using TACACS+, this process occurs:

1. When the connection is established, the switch contacts the TACACS+ daemon to obtain a username prompt, which then appears to the user. The user enters a username, and the switch then contacts the TACACS+ daemon to obtain a password prompt. The switch displays the password prompt to the user, the user enters a password, and the password is then sent to the TACACS+ daemon.

TACACS+ allows a conversation to be held between the daemon and the user until the daemon receives enough information to authenticate the user. The daemon prompts for a username and password combination, but can include other items, such as the user's mother's maiden name.

2. The switch eventually receives one of these responses from the TACACS+ daemon:
 - **ACCEPT**—The user is authenticated and service can begin. If the switch is configured to require authorization, authorization begins at this time.
 - **REJECT**—The user is not authenticated. The user can be denied access or is prompted to retry the login sequence, depending on the TACACS+ daemon.
 - **ERROR**—An error occurred at some time during authentication with the daemon or in the network connection between the daemon and the switch. If an **ERROR** response is received, the switch typically tries to use an alternative method for authenticating the user.
 - **CONTINUE**—The user is prompted for additional authentication information.

After authentication, the user undergoes an additional authorization phase if authorization has been enabled on the switch. Users must first successfully complete TACACS+ authentication before proceeding to TACACS+ authorization.

3. If TACACS+ authorization is required, the TACACS+ daemon is again contacted, and it returns an **ACCEPT** or **REJECT** authorization response. If an **ACCEPT** response is returned, the response contains data in the form of attributes that direct the **EXEC** or **NETWORK** session for that user, determining the services that the user can access:
 - Telnet, Secure Shell (SSH), rlogin, or privileged **EXEC** services
 - Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring TACACS+

This section describes how to configure your switch to support TACACS+. At a minimum, you must identify the host or hosts maintaining the TACACS+ daemon and define the method lists for TACACS+ authentication. You can optionally define method lists for TACACS+ authorization and accounting. A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used, thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

This section contains this configuration information:

- [Default TACACS+ Configuration, page 7-13](#)
- [Identifying the TACACS+ Server Host and Setting the Authentication Key, page 7-13](#)
- [Configuring TACACS+ Login Authentication, page 7-14](#)
- [Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services, page 7-16](#)
- [Starting TACACS+ Accounting, page 7-17](#)

Default TACACS+ Configuration

TACACS+ and AAA are disabled by default.

To prevent a lapse in security, you cannot configure TACACS+ through a network management application. When enabled, TACACS+ can authenticate users accessing the switch through the CLI.

**Note**

Although TACACS+ configuration is performed through the CLI, the TACACS+ server authenticates HTTP connections that have been configured with a privilege level of 15.

Identifying the TACACS+ Server Host and Setting the Authentication Key

You can configure the switch to use a single server or AAA server groups to group existing server hosts for authentication. You can group servers to select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list and contains the list of IP addresses of the selected server hosts.

Beginning in privileged EXEC mode, follow these steps to identify the IP host or host maintaining TACACS+ server and optionally set the encryption key:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	tacacs-server host <i>hostname</i> [port <i>integer</i>] [timeout <i>integer</i>] [key <i>string</i>]	Identify the IP host or hosts maintaining a TACACS+ server. Enter this command multiple times to create a list of preferred hosts. The software searches for hosts in the order in which you specify them. <ul style="list-style-type: none"> For <i>hostname</i>, specify the name or IP address of the host. (Optional) For port <i>integer</i>, specify a server port number. The default is port 49. The range is 1 to 65535. (Optional) For timeout <i>integer</i>, specify a time in seconds the switch waits for a response from the daemon before it times out and declares an error. The default is 5 seconds. The range is 1 to 1000 seconds. (Optional) For key <i>string</i>, specify the encryption key for encrypting and decrypting all traffic between the switch and the TACACS+ daemon. You must configure the same key on the TACACS+ daemon for encryption to be successful.
Step 3	aaa new-model	Enable AAA.
Step 4	aaa group server tacacs+ <i>group-name</i>	(Optional) Define the AAA server-group with a group name. This command puts the switch in a server group subconfiguration mode.
Step 5	server <i>ip-address</i>	(Optional) Associate a particular TACACS+ server with the defined server group. Repeat this step for each TACACS+ server in the AAA server group. Each server in the group must be previously defined in Step 2.
Step 6	end	Return to privileged EXEC mode.
Step 7	show tacacs	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified TACACS+ server name or address, use the **no tacacs-server host** *hostname* global configuration command. To remove a server group from the configuration list, use the **no aaa group server tacacs+** *group-name* global configuration command. To remove the IP address of a TACACS+ server, use the **no server ip-address** server group subconfiguration command.

Configuring TACACS+ Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined. A defined method list overrides the default method list.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2</i> ...]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. For <i>list-name</i>, specify a character string to name the list you are creating. For <i>method1</i>..., specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group tacacs+—Uses TACACS+ authentication. Before you can use this authentication method, you must configure the TACACS+ server. For more information, see the “Identifying the TACACS+ Server Host and Setting the Authentication Key” section on page 7-13. line—Use the line password for authentication. Before you can use this authentication method, you must define a line password. Use the password password line configuration command. local—Use the local username database for authentication. You must enter username information in the database. Use the username password global configuration command. local-case—Use a case-sensitive local username database for authentication. You must enter username information in the database by using the username name password global configuration command. none—Do not use any authentication for login.
Step 4	line [console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.

	Command	Purpose
Step 5	login authentication {default <i>list-name</i> }	Apply the authentication list to a line or set of lines. <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** {default | *list-name*} *method1* [*method2*...] global configuration command. To either disable TACACS+ authentication for logins or to return to the default value, use the **no login authentication** {default | *list-name*} line configuration command.

Configuring TACACS+ Authorization for Privileged EXEC Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is located either in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **tacacs+** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec tacacs+ local** command sets these authorization parameters:

- Use TACACS+ for privileged EXEC access authorization if authentication was performed by using TACACS+.
- Use the local database if authentication was not performed by using TACACS+.



Note

Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify TACACS+ authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network tacacs+	Configure the switch for user TACACS+ authorization for all network-related service requests.
Step 3	aaa authorization exec tacacs+	Configure the switch for user TACACS+ authorization to determine if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Starting TACACS+ Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the TACACS+ security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable TACACS+ accounting for each Cisco IOS privilege level and for network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting network start-stop tacacs+	Enable TACACS+ accounting for all network-related service requests.
Step 3	aaa accounting exec start-stop tacacs+	Enable TACACS+ accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting {network | exec} {start-stop} method1...** global configuration command.

Displaying the TACACS+ Configuration

To display TACACS+ server statistics, use the **show tacacs** privileged EXEC command.

Controlling Switch Access with RADIUS

This section describes how to enable and configure the Remote Authentication Dial-In User Service (RADIUS), which provides detailed accounting information and flexible administrative control over authentication and authorization processes. RADIUS is facilitated through AAA and can be enabled only through AAA commands.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the *Cisco IOS Security Command Reference for Release 12.1*.

This section contains this configuration information:

- [Understanding RADIUS, page 7-18](#)
- [RADIUS Operation, page 7-19](#)
- [Configuring RADIUS, page 7-20](#)
- [Displaying the RADIUS Configuration, page 7-31](#)

Understanding RADIUS

RADIUS is a distributed client/server system that secures networks against unauthorized access. RADIUS clients run on supported Cisco routers and switches. Clients send authentication requests to a central RADIUS server, which contains all user authentication and network service access information. The RADIUS host is normally a multiuser system running RADIUS server software from Cisco (Cisco Secure Access Control Server version 3.0), Livingston, Merit, Microsoft, or another software provider. For more information, refer to the RADIUS server documentation.

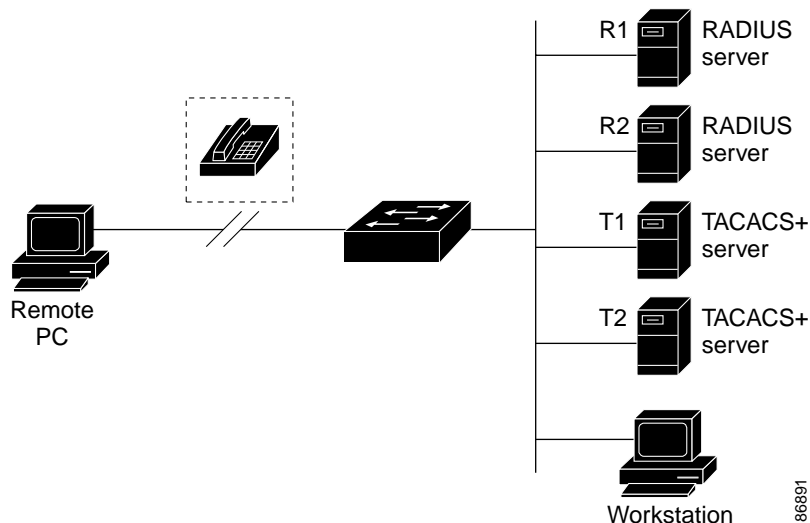
Use RADIUS in these network environments that require access security:

- Networks with multiple-vendor access servers, each supporting RADIUS. For example, access servers from several vendors use a single RADIUS server-based security database. In an IP-based network with multiple vendors' access servers, dial-in users are authenticated through a RADIUS server that has been customized to work with the Kerberos security system.
- Turnkey network security environments in which applications support the RADIUS protocol, such as in an access environment that uses a *smart card* access control system. In one case, RADIUS has been used with Enigma's security cards to validate users and to grant access to network resources.
- Networks already using RADIUS. You can add a Cisco switch containing a RADIUS client to the network. This might be the first step when you make a transition to a TACACS+ server. See [Figure 7-2 on page 7-19](#).
- Network in which the user must only access a single service. Using RADIUS, you can control user access to a single host, to a single utility such as Telnet, or to the network through a protocol such as IEEE 802.1X. For more information about this protocol, see [Chapter 8, "Configuring 802.1X Port-Based Authentication."](#)
- Networks that require resource accounting. You can use RADIUS accounting independently of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, showing the amount of resources (such as time, packets, bytes, and so forth) used during the session. An Internet service provider might use a freeware-based version of RADIUS access control and accounting software to meet special security and billing needs.

RADIUS is not suitable in these network security situations:

- Multiprotocol access environments. RADIUS does not support AppleTalk Remote Access (ARA), NetBIOS Frame Control Protocol (NBFCP), NetWare Asynchronous Services Interface (NASI), or X.25 PAD connections.
- Switch-to-switch or router-to-router situations. RADIUS does not provide two-way authentication. RADIUS can be used to authenticate from one device to a non-Cisco device if the non-Cisco device requires authentication.
- Networks using a variety of services. RADIUS generally binds a user to one service model.

Figure 7-2 Transitioning from RADIUS to TACACS+ Services



RADIUS Operation

When a user attempts to log in and authenticate to a switch that is access controlled by a RADIUS server, these events occur:

1. The user is prompted to enter a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of these responses from the RADIUS server:
 - a. ACCEPT—The user is authenticated.
 - b. REJECT—The user is either not authenticated and is prompted to re-enter the username and password, or access is denied.
 - c. CHALLENGE—A challenge requires additional data from the user.
 - d. CHALLENGE PASSWORD—A response requests the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for privileged EXEC or network authorization. Users must first successfully complete RADIUS authentication before proceeding to RADIUS authorization, if it is enabled. The additional data included with the ACCEPT or REJECT packets includes these items:

- Telnet, SSH, rlogin, or privileged EXEC services
- Connection parameters, including the host or client IP address, access list, and user timeouts

Configuring RADIUS

This section describes how to configure your switch to support RADIUS. At a minimum, you must identify the host or hosts that run the RADIUS server software and define the method lists for RADIUS authentication. You can optionally define method lists for RADIUS authorization and accounting.

A method list defines the sequence and methods to be used to authenticate, to authorize, or to keep accounts on a user. You can use method lists to designate one or more security protocols to be used (such as TACACS+ or local username lookup), thus ensuring a backup system if the initial method fails. The software uses the first method listed to authenticate, to authorize, or to keep accounts on users; if that method does not respond, the software selects the next method in the list. This process continues until there is successful communication with a listed method or the method list is exhausted.

You should have access to and should configure a RADIUS server before configuring RADIUS features on your switch.

This section contains this configuration information:

- [Default RADIUS Configuration, page 7-20](#)
- [Identifying the RADIUS Server Host, page 7-21](#) (required)
- [Configuring RADIUS Login Authentication, page 7-23](#) (required)
- [Defining AAA Server Groups, page 7-25](#) (optional)
- [Configuring RADIUS Authorization for User Privileged Access and Network Services, page 7-27](#) (optional)
- [Starting RADIUS Accounting, page 7-28](#) (optional)
- [Configuring Settings for All RADIUS Servers, page 7-29](#) (optional)
- [Configuring the Switch to Use Vendor-Specific RADIUS Attributes, page 7-29](#) (optional)
- [Configuring the Switch for Vendor-Proprietary RADIUS Server Communication, page 7-31](#) (optional)

Default RADIUS Configuration

RADIUS and AAA are disabled by default.

To prevent a lapse in security, you cannot configure RADIUS through a network management application. When enabled, RADIUS can authenticate users accessing the switch through the CLI.

Identifying the RADIUS Server Host

Switch-to-RADIUS-server communication involves several components:

- Host name or IP address
- Authentication destination port
- Accounting destination port
- Key string
- Timeout period
- Retransmission value

You identify RADIUS security servers by their host name or IP address, host name and specific UDP port numbers, or their IP address and specific UDP port numbers. The combination of the IP address and the UDP port number creates a unique identifier, allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. This unique identifier enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address.

If two different host entries on the same RADIUS server are configured for the same service—for example, accounting—the second host entry configured acts as a fail-over backup to the first one. Using this example, if the first host entry fails to provide accounting services, the switch tries the second host entry configured on the same device for accounting services. (The RADIUS host entries are tried in the order that they are configured.)

A RADIUS server and the switch use a shared secret text string to encrypt passwords and exchange responses. To configure RADIUS to use the AAA security commands, you must specify the host running the RADIUS server daemon and a secret text (key) string that it shares with the switch.

The timeout, retransmission, and encryption key values can be configured globally for all RADIUS servers, on a per-server basis, or in some combination of global and per-server settings. To apply these settings globally to all RADIUS servers communicating with the switch, use the three unique global configuration commands: **radius-server timeout**, **radius-server retransmit**, and **radius-server key**. To apply these values on a specific RADIUS server, use the **radius-server host** global configuration command.



Note

If you configure both global and per-server functions (timeout, retransmission, and key commands) on the switch, the per-server timer, retransmission, and key value commands override global timer, retransmission, and key value commands. For information on configuring these setting on all RADIUS servers, see the [“Configuring Settings for All RADIUS Servers” section on page 7-29](#).

You can configure the switch to use AAA server groups to group existing server hosts for authentication. For more information, see the [“Defining AAA Server Groups” section on page 7-25](#).

Beginning in privileged EXEC mode, follow these steps to configure per-server RADIUS server communication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> • (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. • (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. • (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. • (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. • (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified RADIUS server, use the **no radius-server host** *hostname* | *ip-address* global configuration command.

This example shows how to configure one RADIUS server to be used for authentication and another to be used for accounting:

```
Switch(config)# radius-server host 172.29.36.49 auth-port 1612 key rad1
Switch(config)# radius-server host 172.20.36.50 acct-port 1618 key rad2
```

This example shows how to configure *host1* as the RADIUS server and to use the default ports for both authentication and accounting:

```
Switch(config)# radius-server host host1
```

**Note**

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Configuring RADIUS Login Authentication

To configure AAA authentication, you define a named list of authentication methods and then apply that list to various interfaces. The method list defines the types of authentication to be performed and the sequence in which they are performed; it must be applied to a specific interface before any of the defined authentication methods are performed. The only exception is the default method list (which, by coincidence, is named *default*). The default method list is automatically applied to all interfaces except those that have a named method list explicitly defined.

A method list describes the sequence and authentication methods to be queried to authenticate a user. You can designate one or more security protocols to be used for authentication, thus ensuring a backup system for authentication in case the initial method fails. The software uses the first method listed to authenticate users; if that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access—the authentication process stops, and no other authentication methods are attempted.

Beginning in privileged EXEC mode, follow these steps to configure login authentication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication login { default <i>list-name</i> } <i>method1</i> [<i>method2...</i>]	<p>Create a login authentication method list.</p> <ul style="list-style-type: none"> To create a default list that is used when a named list is <i>not</i> specified in the login authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. For <i>list-name</i>, specify a character string to name the list you are creating. For <i>method1...</i>, specify the actual method the authentication algorithm tries. The additional methods of authentication are used only if the previous method returns an error, not if it fails. <p>Select one of these methods:</p> <ul style="list-style-type: none"> enable—Use the enable password for authentication. Before you can use this authentication method, you must define an enable password by using the enable password global configuration command. group radius—Use RADIUS authentication. Before you can use this authentication method, you must configure the RADIUS server. For more information, see the “Identifying the RADIUS Server Host” section on page 7-21. line—Use the line password for authentication. Before you can use this authentication method, you must define a line password. Use the password password line configuration command. local—Use the local username database for authentication. You must enter username information in the database. Use the username name password global configuration command. local-case—Use a case-sensitive local username database for authentication. You must enter username information in the database by using the username password global configuration command. none—Do not use any authentication for login.
Step 4	line [console tty vty] <i>line-number</i> [<i>ending-line-number</i>]	Enter line configuration mode, and configure the lines to which you want to apply the authentication list.

	Command	Purpose
Step 5	login authentication { default <i>list-name</i> }	Apply the authentication list to a line or set of lines. <ul style="list-style-type: none"> If you specify default, use the default list created with the aaa authentication login command. For <i>list-name</i>, specify the list created with the aaa authentication login command.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable AAA authentication, use the **no aaa authentication login** { **default** | *list-name* } *method1* [*method2*...] global configuration command. To either disable RADIUS authentication for logins or to return to the default value, use the **no login authentication** { **default** | *list-name* } line configuration command.

Defining AAA Server Groups

You can configure the switch to use AAA server groups to group existing server hosts for authentication. You select a subset of the configured server hosts and use them for a particular service. The server group is used with a global server-host list, which lists the IP addresses of the selected server hosts.

Server groups also can include multiple host entries for the same server if each entry has a unique identifier (the combination of the IP address and UDP port number), allowing different ports to be individually defined as RADIUS hosts providing a specific AAA service. If you configure two different host entries on the same RADIUS server for the same service, (for example, accounting), the second configured host entry acts as a fail-over backup to the first one.

You use the **server** group server configuration command to associate a particular server with a defined group server. You can either identify the server by its IP address or identify multiple host instances or entries by using the optional **auth-port** and **acct-port** keywords.

Beginning in privileged EXEC mode, follow these steps to define the AAA server group and associate a particular RADIUS server with it:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } [auth-port <i>port-number</i>] [acct-port <i>port-number</i>] [timeout <i>seconds</i>] [retransmit <i>retries</i>] [key <i>string</i>]	<p>Specify the IP address or host name of the remote RADIUS server host.</p> <ul style="list-style-type: none"> (Optional) For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. (Optional) For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. (Optional) For timeout <i>seconds</i>, specify the time interval that the switch waits for the RADIUS server to reply before resending. The range is 1 to 1000. This setting overrides the radius-server timeout global configuration command setting. If no timeout is set with the radius-server host command, the setting of the radius-server timeout command is used. (Optional) For retransmit <i>retries</i>, specify the number of times a RADIUS request is resent to a server if that server is not responding or responding slowly. The range is 1 to 1000. If no retransmit value is set with the radius-server host command, the setting of the radius-server retransmit global configuration command is used. (Optional) For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. <p>Note The key is a text string that must match the encryption key used on the RADIUS server. Always configure the key as the last item in the radius-server host command. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.</p> <p>To configure the switch to recognize more than one host entry associated with a single IP address, enter this command as many times as necessary, making sure that each UDP port number is different. The switch software searches for hosts in the order in which you specify them. Set the timeout, retransmit, and encryption key values to use with the specific RADIUS host.</p>
Step 3	aaa new-model	Enable AAA.
Step 4	aaa group server radius <i>group-name</i>	<p>Define the AAA server-group with a group name.</p> <p>This command puts the switch in a server group configuration mode.</p>
Step 5	server <i>ip-address</i>	<p>Associate a particular RADIUS server with the defined server group. Repeat this step for each RADIUS server in the AAA server group.</p> <p>Each server in the group must be previously defined in Step 2.</p>
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your entries.

	Command	Purpose
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.
Step 9		Enable RADIUS login authentication. See the “Configuring RADIUS Login Authentication” section on page 7-23.

To remove the specified RADIUS server, use the **no radius-server host** *hostname | ip-address* global configuration command. To remove a server group from the configuration list, use the **no aaa group server radius** *group-name* global configuration command. To remove the IP address of a RADIUS server, use the **no server ip-address** server group configuration command.

In this example, the switch is configured to recognize two different RADIUS group servers (*group1* and *group2*). Group1 has two different host entries on the same RADIUS server configured for the same services. The second host entry acts as a fail-over backup to the first entry.

```
Switch(config)# radius-server host 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config)# radius-server host 172.10.0.1 auth-port 1645 acct-port 1646
Switch(config)# aaa new-model
Switch(config)# aaa group server radius group1
Switch(config-sg-radius)# server 172.20.0.1 auth-port 1000 acct-port 1001
Switch(config-sg-radius)# exit
Switch(config)# aaa group server radius group2
Switch(config-sg-radius)# server 172.20.0.1 auth-port 2000 acct-port 2001
Switch(config-sg-radius)# exit
```

Configuring RADIUS Authorization for User Privileged Access and Network Services

AAA authorization limits the services available to a user. When AAA authorization is enabled, the switch uses information retrieved from the user's profile, which is in the local user database or on the security server, to configure the user's session. The user is granted access to a requested service only if the information in the user profile allows it.

You can use the **aaa authorization** global configuration command with the **radius** keyword to set parameters that restrict a user's network access to privileged EXEC mode.

The **aaa authorization exec radius local** command sets these authorization parameters:

- Use RADIUS for privileged EXEC access authorization if authentication was performed by using RADIUS.
- Use the local database if authentication was not performed by using RADIUS.



Note

Authorization is bypassed for authenticated users who log in through the CLI even if authorization has been configured.

Beginning in privileged EXEC mode, follow these steps to specify RADIUS authorization for privileged EXEC access and network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa authorization network radius	Configure the switch for user RADIUS authorization for all network-related service requests.
Step 3	aaa authorization exec radius	Configure the switch for user RADIUS authorization to determine if the user has privileged EXEC access. The exec keyword might return user profile information (such as autocommand information).
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable authorization, use the **no aaa authorization {network | exec} method1** global configuration command.

Starting RADIUS Accounting

The AAA accounting feature tracks the services that users are accessing and the amount of network resources that they are consuming. When AAA accounting is enabled, the switch reports user activity to the RADIUS security server in the form of accounting records. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can then be analyzed for network management, client billing, or auditing.

Beginning in privileged EXEC mode, follow these steps to enable RADIUS accounting for each Cisco IOS privilege level and for network services:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa accounting network start-stop radius	Enable RADIUS accounting for all network-related service requests.
Step 3	aaa accounting exec start-stop radius	Enable RADIUS accounting to send a start-record accounting notice at the beginning of a privileged EXEC process and a stop-record at the end.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable accounting, use the **no aaa accounting {network | exec} {start-stop} method1...** global configuration command.

Configuring Settings for All RADIUS Servers

Beginning in privileged EXEC mode, follow these steps to configure global communication settings between the switch and all RADIUS servers:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and all RADIUS servers. Note The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 3	radius-server retransmit <i>retries</i>	Specify the number of times the switch sends each RADIUS request to the server before giving up. The default is 3; the range 1 to 1000.
Step 4	radius-server timeout <i>seconds</i>	Specify the number of seconds a switch waits for a reply to a RADIUS request before resending the request. The default is 5 seconds; the range is 1 to 1000.
Step 5	radius-server deadtime <i>minutes</i>	Specify the number of minutes a RADIUS server, which is not responding to authentication requests, to be skipped, thus avoiding the wait for the request to timeout before trying the next configured server. The default is 0; the range is 1 to 1440 minutes.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config	Verify your settings.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting for the retransmit, timeout, and deadtime, use the **no** forms of these commands.

Configuring the Switch to Use Vendor-Specific RADIUS Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific information between the switch and the RADIUS server by using the vendor-specific attribute (attribute 26). Vendor-specific attributes (VSAs) allow vendors to support their own extended attributes not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option by using the format recommended in the specification. Cisco's vendor-ID is 9, and the supported option has vendor-type 1, which is named *cisco-avpair*. The value is a string with this format:

```
protocol : attribute sep value *
```

Protocol is a value of the Cisco protocol attribute for a particular type of authorization. *Attribute* and *value* are an appropriate attribute-value (AV) pair defined in the Cisco TACACS+ specification, and *sep* is = for mandatory attributes and is * for optional attributes. The full set of features available for TACACS+ authorization can then be used for RADIUS.

For example, this AV pair activates Cisco's *multiple named ip address pools* feature during IP authorization (during PPP IPCP address assignment):

```
cisco-avpair= "ip:addr-pool=first"
```

This example shows how to provide a user logging in from a switch with immediate access to privileged EXEC commands:

```
cisco-avpair= "shell:priv-lvl=15"
```

This example shows how to specify an authorized VLAN in the RADIUS server database:

```
cisco-avpair= "tunnel-type(#64)=VLAN(13)"
cisco-avpair= "tunnel-medium-type(#65)=802 media(6)"
cisco-avpair= "tunnel-private-group-ID(#81)=vlanid"
```

This example shows how to apply an input ACL in ASCII format to an interface for the duration of this connection:

```
cisco-avpair= "ip:inacl#1=deny ip 10.10.10.10 0.0.255.255 20.20.20.20 255.255.0.0"
cisco-avpair= "ip:inacl#2=deny ip 10.10.10.10 0.0.255.255 any"
cisco-avpair= "mac:inacl#3=deny any any decnet-iv"
```

This example shows how to apply an output ACL in ASCII format to an interface for the duration of this connection:

```
cisco-avpair= "ip:outacl#2=deny ip 10.10.10.10 0.0.255.255 any"
```

Other vendors have their own unique vendor-IDs, options, and associated VSAs. For more information about vendor-IDs and VSAs, refer to RFC 2138, "Remote Authentication Dial-In User Service (RADIUS)."

Beginning in privileged EXEC mode, follow these steps to configure the switch to recognize and use VSAs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server vsa send [accounting authentication]	<p>Enable the switch to recognize and use VSAs as defined by RADIUS IETF attribute 26.</p> <ul style="list-style-type: none"> (Optional) Use the accounting keyword to limit the set of recognized vendor-specific attributes to only accounting attributes. (Optional) Use the authentication keyword to limit the set of recognized vendor-specific attributes to only authentication attributes. <p>If you enter this command without keywords, both accounting and authentication vendor-specific attributes are used.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your settings.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For a complete list of RADIUS attributes or more information about vendor-specific attribute 26, refer to the "RADIUS Attributes" appendix in the *Cisco IOS Security Configuration Guide for Release 12.1*.

Configuring the Switch for Vendor-Proprietary RADIUS Server Communication

Although an IETF draft standard for RADIUS specifies a method for communicating vendor-proprietary information between the switch and the RADIUS server, some vendors have extended the RADIUS attribute set in a unique way. Cisco IOS software supports a subset of vendor-proprietary RADIUS attributes.

As mentioned earlier, to configure RADIUS (whether vendor-proprietary or IETF draft-compliant), you must specify the host running the RADIUS server daemon and the secret text string it shares with the switch. You specify the RADIUS host and secret text string by using the **radius-server** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to specify a vendor-proprietary RADIUS server host and a shared secret text string:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } non-standard	Specify the IP address or host name of the remote RADIUS server host and identify that it is using a vendor-proprietary implementation of RADIUS.
Step 3	radius-server key <i>string</i>	Specify the shared secret text string used between the switch and the vendor-proprietary RADIUS server. The switch and the RADIUS server use this text string to encrypt passwords and exchange responses. Note The key is a text string that must match the encryption key used on the RADIUS server. Leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in your key, do not enclose the key in quotation marks unless the quotation marks are part of the key.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the vendor-proprietary RADIUS host, use the **no radius-server host** {*hostname* | *ip-address*} **non-standard** global configuration command. To disable the key, use the **no radius-server key** global configuration command.

This example shows how to specify a vendor-proprietary RADIUS host and to use a secret key of *rad124* between the switch and the server:

```
Switch(config)# radius-server host 172.20.30.15 nonstandard
Switch(config)# radius-server key rad124
```

Displaying the RADIUS Configuration

To display the RADIUS configuration, use the **show running-config** privileged EXEC command.

Controlling Switch Access with Kerberos

This section describes how to enable and configure the Kerberos security system, which authenticates requests for network resources by using a trusted third party. To use this feature, the cryptographic (that is, supports encryption) version of the switch software must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, refer to the release notes for this release.

This section consists of these topics:

- [Understanding Kerberos, page 7-32](#)
- [Kerberos Operation, page 7-34](#)
- [Configuring Kerberos, page 7-36](#)

For Kerberos configuration examples, refer to the “Kerberos Configuration Examples” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgr/secur_c/scprt2/.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the “Kerberos Commands” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Command Reference, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgr/secur_r/srprt2/srdkerb.htm.

**Note**

In the Kerberos configuration examples and in the *Cisco IOS Security Command Reference, Release 12.1*, the trusted third party can be a Catalyst 2970 switch that supports Kerberos, that is configured as a network security server, and that can authenticate users by using the Kerberos protocol.

Understanding Kerberos

Kerberos is a secret-key network authentication protocol, which was developed at the Massachusetts Institute of Technology (MIT). It uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication and authenticates requests for network resources. Kerberos uses the concept of a trusted third party to perform secure verification of users and services. This trusted third party is called the *key distribution center* (KDC).

Kerberos verifies that users are who they claim to be and the network services that they use are what the services claim to be. To do this, a KDC or trusted Kerberos server issues tickets to users. These tickets, which have a limited lifespan, are stored in user credential caches. The Kerberos server uses the tickets instead of usernames and passwords to authenticate users and network services.

**Note**

A Kerberos server can be a Catalyst 2970 switch that is configured as a network security server and that can authenticate users by using the Kerberos protocol.

The Kerberos credential scheme uses a process called *single login*. This process authenticates a user once and then allows secure authentication (without encrypting another password) wherever that user credential is accepted.

This software release supports Kerberos 5, which allows organizations that are already using Kerberos 5 to use the same Kerberos authentication database on the KDC that they are already using on their other network hosts (such as UNIX servers and PCs).

In this software release, Kerberos supports these network services:

- Telnet
- rlogin
- rsh (Remote Shell Protocol)

Table 7-2 lists the common Kerberos-related terms and definitions:

Table 7-2 Kerberos Terms

Term	Definition
Authentication	A process by which a user or service identifies itself to another service. For example, a client can authenticate to a switch or a switch can authenticate to another switch.
Authorization	A means by which the switch determines what privileges the user has in a network or on the switch and what actions the user can perform.
Credential	A general term that refers to authentication tickets, such as TGTs ¹ and service credentials. Kerberos credentials verify the identity of a user or service. If a network service decides to trust the Kerberos server that issued a ticket, it can be used in place of re-entering a username and password. Credentials have a default lifespan of eight hours.
Instance	<p>An authorization level label for Kerberos principals. Most Kerberos principals are of the form <i>user@REALM</i> (for example, <i>smith@EXAMPLE.COM</i>). A Kerberos principal with a Kerberos instance has the form <i>user/instance@REALM</i> (for example, <i>smith/admin@EXAMPLE.COM</i>). The Kerberos instance can be used to specify the authorization level for the user if authentication is successful. The server of each network service might implement and enforce the authorization mappings of Kerberos instances but is not required to do so.</p> <p>Note The Kerberos principal and instance names <i>must</i> be in all lowercase characters.</p> <p>Note The Kerberos realm name <i>must</i> be in all uppercase characters.</p>
KDC ²	Key distribution center that consists of a Kerberos server and database program that is running on a network host.
Kerberized	A term that describes applications and services that have been modified to support the Kerberos credential infrastructure.
Kerberos realm	<p>A domain consisting of users, hosts, and network services that are registered to a Kerberos server. The Kerberos server is trusted to verify the identity of a user or network service to another user or network service.</p> <p>Note The Kerberos realm name <i>must</i> be in all uppercase characters.</p>
Kerberos server	A daemon that is running on a network host. Users and network services register their identity with the Kerberos server. Network services query the Kerberos server to authenticate to other network services.

Table 7-2 Kerberos Terms (continued)

Term	Definition
KEYTAB ³	A password that a network service shares with the KDC. In Kerberos 5 and later Kerberos versions, the network service authenticates an encrypted service credential by using the KEYTAB to decrypt it. In Kerberos versions earlier than Kerberos 5, KEYTAB is referred to as SRVTAB ⁴ .
Principal	Also known as a Kerberos identity, this is who you are or what a service is according to the Kerberos server. Note The Kerberos principal name <i>must</i> be in all lowercase characters.
Service credential	A credential for a network service. When issued from the KDC, this credential is encrypted with the password shared by the network service and the KDC. The password is also shared with the user TGT.
SRVTAB	A password that a network service shares with the KDC. In Kerberos 5 or later Kerberos versions, SRVTAB is referred to as KEYTAB.
TGT	Ticket granting ticket that is a credential that the KDC issues to authenticated users. When users receive a TGT, they can authenticate to network services within the Kerberos realm represented by the KDC.

1. TGT = ticket granting ticket
2. KDC = key distribution center
3. KEYTAB = key table
4. SRVTAB = server table

Kerberos Operation

A Kerberos server can be a Catalyst 2970 switch that is configured as a network security server and that can authenticate remote users by using the Kerberos protocol. Although you can customize Kerberos in a number of ways, remote users attempting to access network services must pass through three layers of security before they can access network services.

To authenticate to network services by using a Catalyst 2970 switch as a Kerberos server, remote users must follow these steps:

1. [Authenticating to a Boundary Switch, page 7-35](#)
2. [Obtaining a TGT from a KDC, page 7-35](#)
3. [Authenticating to Network Services, page 7-35](#)

Authenticating to a Boundary Switch

This section describes the first layer of security through which a remote user must pass. The user must first authenticate to the boundary switch. This process then occurs:

1. The user opens an un-Kerberized Telnet connection to the boundary switch.
2. The switch prompts the user for a username and password.
3. The switch requests a TGT from the KDC for this user.
4. The KDC sends an encrypted TGT that includes the user identity to the switch.
5. The switch attempts to decrypt the TGT by using the password that the user entered.
 - If the decryption is successful, the user is authenticated to the switch.
 - If the decryption is not successful, the user repeats Step 2 either by re-entering the username and password (noting if Caps Lock or Num Lock is on or off) or by entering a different username and password.

A remote user who initiates a un-Kerberized Telnet session and authenticates to a boundary switch is inside the firewall, but the user must still authenticate directly to the KDC before getting access to the network services. The user must authenticate to the KDC because the TGT that the KDC issues is stored on the switch and cannot be used for additional authentication until the user logs on to the switch.

Obtaining a TGT from a KDC

This section describes the second layer of security through which a remote user must pass. The user must now authenticate to a KDC and obtain a TGT from the KDC to access network services.

For instructions about how to authenticate to a KDC, refer to the “Obtaining a TGT from a KDC” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/scdkerb.htm#xtocid154005.

Authenticating to Network Services

This section describes the third layer of security through which a remote user must pass. The user with a TGT must now authenticate to the network services in a Kerberos realm.

For instructions about how to authenticate to a network service, refer to the “Authenticating to Network Services” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt2/scdkerb.htm#xtocid154006.

Configuring Kerberos

So that remote users can authenticate to network services, you must configure the hosts and the KDC in the Kerberos realm to communicate and mutually authenticate users and network services. To do this, you must identify them to each other. You add entries for the hosts to the Kerberos database on the KDC and add KEYTAB files generated by the KDC to all hosts in the Kerberos realm. You also create entries for the users in the KDC database.

When you add or create entries for the hosts and users, follow these guidelines:

- The Kerberos principal name *must* be in all lowercase characters.
- The Kerberos instance name *must* be in all lowercase characters.
- The Kerberos realm name *must* be in all uppercase characters.



Note

A Kerberos server can be a Catalyst 2970 switch that is configured as a network security server and that can authenticate users by using the Kerberos protocol.

To set up a Kerberos-authenticated server-client system, follow these steps:

- Configure the KDC by using Kerberos commands.
- Configure the switch to use the Kerberos protocol.

For instructions, refer to the “Kerberos Configuration Task List” section in the “Security Server Protocols” chapter of the *Cisco IOS Security Configuration Guide, Release 12.1*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgr/secur_c/scprt2/scdkerb.htm#xtocid154007.

Configuring the Switch for Local Authentication and Authorization

You can configure AAA to operate without a server by setting the switch to implement AAA in local mode. The switch then handles authentication and authorization. No accounting is available in this configuration.

Beginning in privileged EXEC mode, follow these steps to configure the switch for local AAA:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication login default local	Set the login authentication to use the local username database. The default keyword applies the local user database authentication to all interfaces.
Step 4	aaa authorization exec local	Configure user AAA authorization to determine if the user is allowed to run an EXEC shell by checking the local database.
Step 5	aaa authorization network local	Configure user AAA authorization for all network-related service requests.

	Command	Purpose
Step 6	username <i>name</i> [privilege <i>level</i>] { password <i>encryption-type password</i> }	Enter the local database, and establish a username-based authentication system. Repeat this command for each user. <ul style="list-style-type: none"> For <i>name</i>, specify the user ID as one word. Spaces and quotation marks are not allowed. (Optional) For <i>level</i>, specify the privilege level the user has after gaining access. The range is 0 to 15. Level 15 gives privileged EXEC mode access. Level 0 gives user EXEC mode access. For <i>encryption-type</i>, enter 0 to specify that an unencrypted password follows. Enter 7 to specify that a hidden password follows. For <i>password</i>, specify the password the user must enter to gain access to the switch. The password must be from 1 to 25 characters, can contain embedded spaces, and must be the last option specified in the username command.
Step 7	end	Return to privileged EXEC mode.
Step 8	show running-config	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable authorization, use the **no aaa authorization** {**network** | **exec**} *method1* global configuration command.

Configuring the Switch for Secure Shell

This section describes how to configure the Secure Shell (SSH) feature. To use this feature, the cryptographic (encrypted) software image must be installed on your switch. You must obtain authorization to use this feature and to download the cryptographic software files from Cisco.com. For more information, refer to the release notes for this release.

This section contains this information:

- [Understanding SSH, page 7-38](#)
- [Configuring SSH, page 7-39](#)
- [Displaying the SSH Configuration and Status, page 7-41](#)

For SSH configuration examples, refer to the “SSH Configuration Examples” section in the “Configuring Secure Shell” chapter of the *Cisco IOS Security Configuration Guide, Cisco IOS Release 12.2*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/fothersf/scfssh.htm



Note

For complete syntax and usage information for the commands used in this section, refer to the command reference for this release and the command reference for Cisco IOS Release 12.2 at this URL:

<http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/index.htm>.

Understanding SSH

SSH is a protocol that provides a secure, remote connection to a device. SSH provides more security for remote connections than Telnet does by providing strong encryption when a device is authenticated. This software release supports SSH version 1 (SSHv1) and SSH version 2 (SSHv2).

This section consists of these topics:

- [SSH Servers, Integrated Clients, and Supported Versions, page 7-38](#)
- [Limitations, page 7-38](#)

SSH Servers, Integrated Clients, and Supported Versions

The SSH feature has an SSH server and an SSH integrated client, which are applications that run on the switch. You can use an SSH client to connect to a switch running the SSH server. The SSH server works with the SSH client supported in this release and with non-Cisco SSH clients. The SSH client also works with the SSH server supported in this release and with non-Cisco SSH servers.

The switch supports an SSHv1 or an SSHv2 server.

The switch supports an SSHv1 client.

SSH supports the Data Encryption Standard (DES) encryption algorithm, the Triple DES (3DES) encryption algorithm, and password-based user authentication.

SSH also supports these user authentication methods:

- TACACS+ (for more information, see the [“Controlling Switch Access with TACACS+”](#) section on [page 7-10](#))
- RADIUS (for more information, see the [“Controlling Switch Access with RADIUS”](#) section on [page 7-18](#))
- Local authentication and authorization (for more information, see the [“Configuring the Switch for Local Authentication and Authorization”](#) section on [page 7-36](#))



Note

This software release does not support IP Security (IPSec).

Limitations

These limitations apply to SSH:

- The switch supports Rivest, Shamir, and Adelman (RSA) authentication.
- SSH supports only the execution-shell application.

- The SSH server and the SSH client are supported only on DES (56-bit) and 3DES (168-bit) data encryption software.
- The switch does not support the Advanced Encryption Standard (AES) symmetric encryption algorithm.

Configuring SSH

This section has this configuration information:

- [Configuration Guidelines, page 7-39](#)
- [Setting Up the Switch to Run SSH, page 7-39](#) (required)
- [Configuring the SSH Server, page 7-40](#) (required only if you are configuring the switch as an SSH server)

Configuration Guidelines

Follow these guidelines when configuring the switch as an SSH server or SSH client:

- An RSA key pair generated by a SSHv1 server can be used by an SSHv2 server, and the reverse.
- If you get CLI error messages after entering the **crypto key generate rsa** global configuration command, an RSA key pair has not been generated. Reconfigure the host name and domain, and then enter the **crypto key generate rsa** command. For more information, see the [“Setting Up the Switch to Run SSH” section on page 7-39](#).
- When generating the RSA key pair, the message `No host name specified` might appear. If it does, you must configure a host name by using the **hostname** global configuration command.
- When generating the RSA key pair, the message “No domain specified” might appear. If it does, you must configure an IP domain name by using the **ip domain-name** global configuration command.
- When configuring the local authentication and authorization authentication method, make sure that AAA is disabled on the console.

Setting Up the Switch to Run SSH

Follow these steps to set up your switch to run SSH:

1. Download the cryptographic software image from Cisco.com. This step is required. For more information, refer to the release notes for this release.
2. Configure a host name and IP domain name for the switch. Follow this procedure only if you are configuring the switch as an SSH server.
3. Generate an RSA key pair for the switch, which automatically enables SSH. Follow this procedure only if you are configuring the switch as an SSH server.
4. Configure user authentication for local or remote access. This step is required. For more information, see the [“Configuring the Switch for Local Authentication and Authorization” section on page 7-36](#).

Beginning in privileged EXEC mode, follow these steps to configure a host name and an IP domain name and to generate an RSA key pair. This procedure is required if you are configuring the switch as an SSH server.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	hostname <i>hostname</i>	Configure a host name for your switch.
Step 3	ip domain-name <i>domain_name</i>	Configure a host domain for your switch.
Step 4	crypto key generate rsa	<p>Enable the SSH server for local and remote authentication on the switch and generate an RSA key pair.</p> <p>We recommend that a minimum modulus size of 1024 bits.</p> <p>When you generate RSA keys, you are prompted to enter a modulus length. A longer modulus length might be more secure, but it takes longer to generate and to use.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip ssh or show ssh	<p>Show the version and configuration information for your SSH server.</p> <p>Show the status of the SSH server on the switch.</p>
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the RSA key pair, use the **crypto key zeroize rsa** global configuration command. After the RSA key pair is deleted, the SSH server is automatically disabled.

Configuring the SSH Server

Beginning in privileged EXEC mode, follow these steps to configure the SSH server:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip ssh version [1 2]	<p>(Optional) Configure the switch to run SSH version 1 or SSH version 2.</p> <ul style="list-style-type: none"> 1—Configure the switch to run SSH version 1. 2—Configure the switch to run SSH version 2. <p>If you do not enter this command or do not specify a keyword, the SSH server selects the latest SSH version supported by the SSH client. For example, if the SSH client supports SSHv1 and SSHv2, the SSH server selects SSHv2.</p>

	Command	Purpose
Step 3	ip ssh {timeout seconds authentication-retries number}	<p>Configure the SSH control parameters:</p> <ul style="list-style-type: none"> Specify the time-out value in seconds; the default is 120 seconds. The range is 0 to 120 seconds. This parameter applies to the SSH negotiation phase. After the connection is established, the switch uses the default time-out values of the CLI-based sessions. <p>By default, up to five simultaneous, encrypted SSH connections for multiple CLI-based sessions over the network are available (session 0 to session 4). After the execution shell starts, the CLI-based session time-out value returns to the default of 10 minutes.</p> <ul style="list-style-type: none"> Specify the number of times that a client can re-authenticate to the server. The default is 3; the range is 0 to 5. <p>Repeat this step when configuring both parameters.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip ssh or show ssh	<p>Show the version and configuration information for your SSH server.</p> <p>Show the status of the SSH server connections on the switch.</p>
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default SSH control parameters, use the **no ip ssh {timeout | authentication-retries}** global configuration command.

Displaying the SSH Configuration and Status

To display the SSH server configuration and status, use one or more of the privileged EXEC commands in [Table 7-3](#):

Table 7-3 Commands for Displaying the SSH Server Configuration and Status

Command	Purpose
show ip ssh	Shows the version and configuration information for the SSH server.
show ssh	Shows the status of the SSH server.

For more information about these commands, refer to the “Secure Shell Commands” section in the “Other Security Features” chapter of the *Cisco IOS Security Command Reference, Cisco IOS Release 12.2*, at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/fothercr/srfssh.htm.



Configuring 802.1X Port-Based Authentication

This chapter describes how to configure IEEE 802.1X port-based authentication on the Catalyst 2970 switch. As LANs extend to hotels, airports, and corporate lobbies, creating insecure environments, 802.1X prevents unauthorized devices (clients) from gaining access to the network.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding 802.1X Port-Based Authentication, page 8-1](#)
- [Configuring 802.1X Authentication, page 8-9](#)
- [Displaying 802.1X Statistics and Status, page 8-19](#)

Understanding 802.1X Port-Based Authentication

The IEEE 802.1X standard defines a client-server-based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. The authentication server authenticates each client connected to a switch port before making available any services offered by the switch or the LAN.

Until the client is authenticated, 802.1X access control allows only Extensible Authentication Protocol over LAN (EAPOL), Cisco Discovery Protocol (CDP), and Spanning Tree Protocol (STP) traffic through the port to which the client is connected. After authentication is successful, normal traffic can pass through the port.

These sections describe 802.1X port-based authentication:

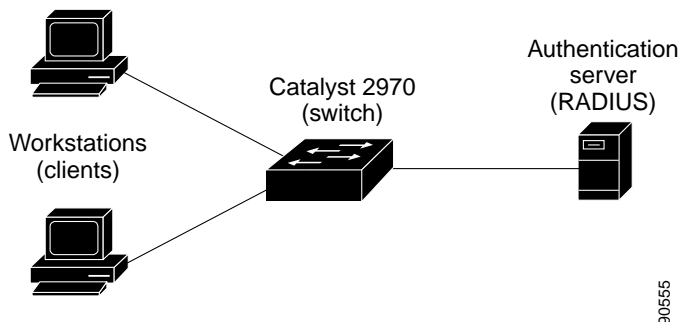
- [Device Roles, page 8-2](#)
- [Authentication Initiation and Message Exchange, page 8-3](#)
- [Ports in Authorized and Unauthorized States, page 8-4](#)
- [Supported Topologies, page 8-5](#)
- [Using 802.1X with Port Security, page 8-6](#)
- [Using 802.1X with Voice VLAN Ports, page 8-6](#)
- [Using 802.1X with VLAN Assignment, page 8-7](#)
- [Using 802.1X with Guest VLAN, page 8-8](#)

- [Using 802.1X with Per-User ACLs, page 8-9](#)

Device Roles

With 802.1X port-based authentication, the devices in the network have specific roles as shown in [Figure 8-1](#).

Figure 8-1 802.1X Device Roles



- *Client*—the device (workstation) that requests access to the LAN and switch services and responds to requests from the switch. The workstation must be running 802.1X-compliant client software such as that offered in the Microsoft Windows XP operating system. (The client is the *supplicant* in the IEEE 802.1X specification.)



Note

To resolve Windows XP network connectivity and 802.1X authentication issues, read the Microsoft Knowledge Base article at this URL:

<http://support.microsoft.com/support/kb/articles/Q303/5/97.ASP>

- *Authentication server*—performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch whether or not the client is authorized to access the LAN and switch services. Because the switch acts as the proxy, the authentication service is transparent to the client. In this release, the Remote Authentication Dial-In User Service (RADIUS) security system with Extensible Authentication Protocol (EAP) extensions is the only supported authentication server. It is available in Cisco Secure Access Control Server version 3.0 or later. RADIUS operates in a client/server model in which secure authentication information is exchanged between the RADIUS server and one or more RADIUS clients.
- *Switch* (edge switch or wireless access point)—controls the physical access to the network based on the authentication status of the client. The switch acts as an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch includes the RADIUS client, which is responsible for encapsulating and decapsulating the EAP frames and interacting with the authentication server.

When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is re-encapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

The devices that can act as intermediaries include the Catalyst 3750, Catalyst 3550, Catalyst 2970, Catalyst 2955, Catalyst 2950, Catalyst 2940 switches, or a wireless access point. These devices must be running software that supports the RADIUS client and 802.1X.

Authentication Initiation and Message Exchange

The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link state transitions from down to up. It then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during bootup, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.



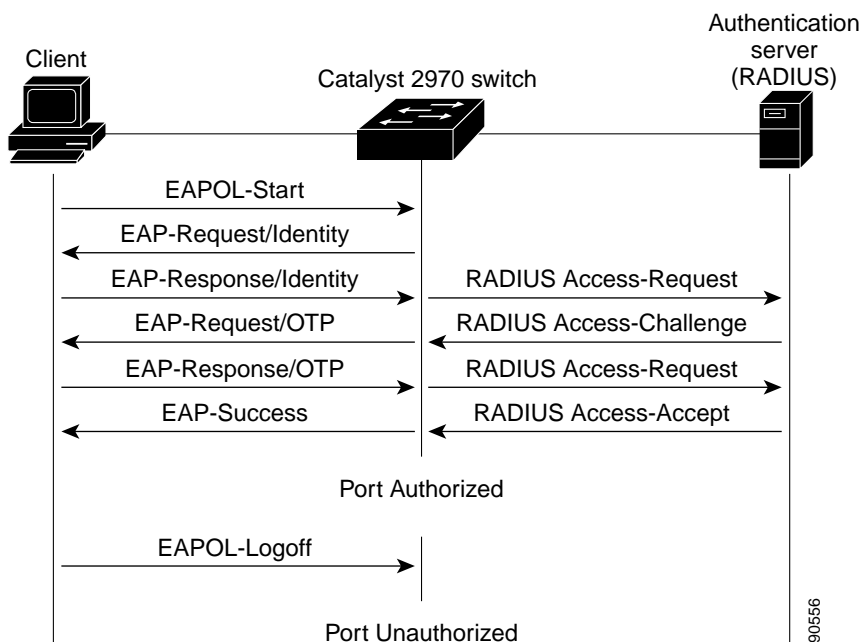
Note

If 802.1X is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client sends frames as if the port is in the authorized state. A port in the authorized state effectively means that the client has been successfully authenticated. For more information, see the [“Ports in Authorized and Unauthorized States” section on page 8-4](#).

When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized. For more information, see the [“Ports in Authorized and Unauthorized States” section on page 8-4](#).

The specific exchange of EAP frames depends on the authentication method being used. [Figure 8-2](#) shows a message exchange initiated by the client using the One-Time-Password (OTP) authentication method with a RADIUS server.

Figure 8-2 Message Exchange



Ports in Authorized and Unauthorized States

The switch port state determines whether or not the client is granted access to the network. The port starts in the *unauthorized* state. While in this state, the port disallows all ingress and egress traffic except for 802.1X, CDP, and STP protocol packets. When a client is successfully authenticated, the port transitions to the *authorized* state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1X is connected to an unauthorized 802.1X port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network.

In contrast, when an 802.1X-enabled client connects to a port that is not running the 802.1X protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request for a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You control the port authorization state by using the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized**—disables 802.1X authentication and causes the port to transition to the authorized state without any authentication exchange required. The port sends and receives normal traffic without 802.1X-based authentication of the client. This is the default setting.
- **force-unauthorized**—causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.
- **auto**—enables 802.1X authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is

received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. Each client attempting to access the network is uniquely identified by the switch by using the client MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If the authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can resend the request. If no response is received from the server after the specified number of attempts, authentication fails, and network access is not granted.

When a client logs off, it sends an EAPOL-logoff message, causing the switch port to transition to the unauthorized state.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received, the port returns to the unauthorized state.

Supported Topologies

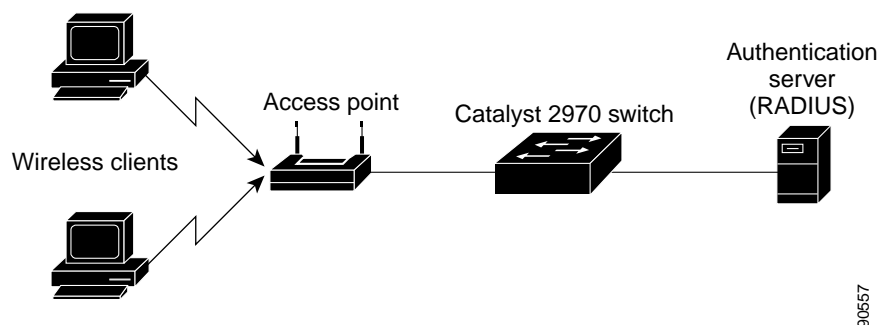
The 802.1X port-based authentication is supported in two topologies:

- Point-to-point
- Wireless LAN

In a point-to-point configuration (see [Figure 8-1 on page 8-2](#)), only one client can be connected to the 802.1X-enabled switch port. The switch detects the client when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

[Figure 8-3](#) shows 802.1X port-based authentication in a wireless LAN. The 802.1X port is configured as a multiple-hosts port that becomes authorized as soon as one client is authenticated. When the port is authorized, all other hosts indirectly attached to the port are granted access to the network. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), the switch denies access to the network to all of the attached clients. In this topology, the wireless access point is responsible for authenticating the clients attached to it, and the wireless access point acts as a client to the switch.

Figure 8-3 Wireless LAN Example



90557

Using 802.1X with Port Security

You can configure 802.1X port and port security in either single-host or multiple-hosts mode. (You also must configure port security on the port by using the **switchport port-security** interface configuration command.) When you enable port security and 802.1X on a port, 802.1X authenticates the port, and port security manages network access for all MAC addresses, including that of the client. You can then limit the number or group of clients that can access the network through an 802.1X port.

These are some examples of the interaction between 802.1X and port security on the switch:

- When a client is authenticated, and the port security table is not full, the client MAC address is added to the port security list of secure hosts. The port then proceeds to come up normally.

When a client is authenticated and manually configured for port security, it is guaranteed an entry in the secure host table (unless port security static aging has been enabled).

A security violation occurs if the client is authenticated, but the port security table is full. This can happen if the maximum number of secure hosts has been statically configured or if the client ages out of the secure host table. If the client address is aged, its place in the secure host table can be taken by another host.

If the security violation is caused by the first authenticated host, the interface becomes error-disabled and immediately shuts down.

The port security violation modes determine the action for security violations. For more information, see the [“Security Violations” section on page 19-8](#).

- When you manually remove an 802.1X client address from the port security table by using the **no switchport port-security mac-address mac-address** interface configuration command, you should re-authenticate the 802.1X client by using the **dot1x re-authenticate interface interface-id** privileged EXEC command.
- When an 802.1X client logs off, the port transitions to an unauthenticated state, and all dynamic entries in the secure host table are cleared, including the entry for the client. Normal authentication then takes place.
- If the port is administratively shut down, the port becomes unauthenticated, and all dynamic entries are removed from the secure host table.
- Port security and a voice VLAN can be configured simultaneously on an 802.1X port that is in either single-host or multiple-hosts mode. Port security applies to both the voice VLAN identifier (VVID) and the port VLAN identifier (PVID).

For more information about enabling port security on your switch, see the [“Configuring Port Security” section on page 19-7](#).

Using 802.1X with Voice VLAN Ports

A voice VLAN port is a special access port associated with two VLAN identifiers:

- VVID to carry voice traffic to and from the IP phone. The VVID is used to configure the IP phone connected to the port.
- PVID to carry the data traffic to and from the workstation connected to the switch through the IP phone. The PVID is the native VLAN of the port.

Each port that you configure for a voice VLAN is associated with a PVID and a VVID. This configuration allows voice traffic and data traffic to be separated onto different VLANs.

Before Cisco IOS Release 12.1(14)EA1, a switch in single-host mode accepted traffic from a single host, and voice traffic was not allowed. In multiple-hosts mode, the switch did not accept voice traffic until the client was authenticated on the primary VLAN, thus making the IP phone inoperable until the user logged in.

With Cisco IOS Release 12.1(14)EA1 and later, the IP phone uses the VVID for its voice traffic regardless of the authorized or unauthorized state of the port. This allows the phone to work independently of 802.1X authentication.

When you enable the single-host mode, multiple IP phones are allowed on the VVID; only one 802.1X client is allowed on the PVID. When you enable the multiple-hosts mode and when an 802.1X user is authenticated on the primary VLAN, additional clients on the voice VLAN are unrestricted after 802.1X authentication succeeds on the primary VLAN.

A voice VLAN port becomes active when there is link, and the device MAC address appears after the first CDP message from the IP phone. Cisco IP phones do not relay CDP messages from other devices. As a result, if several IP phones are connected in series, the switch recognizes only the one directly connected to it. When 802.1X is enabled on a voice VLAN port, the switch drops packets from unrecognized IP phones more than one hop away.

When 802.1X is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.

For more information about voice VLANs, see the [Chapter 13, “Configuring Voice VLAN.”](#)

Using 802.1X with VLAN Assignment

Before Cisco IOS Release 12.1(14)EA1, when an 802.1X port was authenticated, it was authorized to be in the access VLAN configured on the port even if the RADIUS server returned an authorized VLAN from its database. Recall that an access VLAN is a VLAN assigned to an access port. All packets sent from or received on this port belong to this VLAN.

However, with Cisco IOS Release 12.1(14)EA1 and later, the switch supports 802.1X with VLAN assignment. After successful 802.1X authentication of a port, the RADIUS server sends the VLAN assignment to configure the switch port. The RADIUS server database maintains the username-to-VLAN mappings, which assigns the VLAN based on the username of the client connected to the switch port. You can use this feature to limit network access for certain users.

When configured on the switch and the RADIUS server, 802.1X with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server or if 802.1X authorization is disabled, the port is configured in its access VLAN after successful authentication.
- If 802.1X authorization is enabled but the VLAN information from the RADIUS server is not valid, the port returns to the unauthorized state and remains in the configured access VLAN. This prevents ports from appearing unexpectedly in an inappropriate VLAN because of a configuration error.
Configuration errors could include specifying a malformed VLAN ID, a nonexistent VLAN ID, or an attempted assignment to a voice VLAN ID.
- If 802.1X authorization is enabled and all information from the RADIUS server is valid, the port is placed in the specified VLAN after authentication.
- If the multiple-hosts mode is enabled on an 802.1X port, all hosts are placed in the same VLAN (specified by the RADIUS server) as the first authenticated host.
- If 802.1X and port security are enabled on a port, the port is placed in RADIUS server assigned VLAN.
- If 802.1X is disabled on the port, it is returned to the configured access VLAN.

When the port is in the force authorized, force unauthorized, unauthorized, or shutdown state, it is put into the configured access VLAN.

If an 802.1X port is authenticated and put in the RADIUS server assigned VLAN, any change to the port access VLAN configuration does not take effect.

The 802.1X with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VLAN Membership Policy Server (VMPS).

To configure VLAN assignment you need to perform these tasks:

- Enable AAA authorization by using the **network** keyword to allow interface configuration from the RADIUS server.
- Enable 802.1X. (The VLAN assignment feature is automatically enabled when you configure 802.1X on an access port).
- Assign vendor-specific tunnel attributes in the RADIUS server. The RADIUS server must return these attributes to the switch:
 - [64] Tunnel-Type = VLAN
 - [65] Tunnel-Medium-Type = 802
 - [81] Tunnel-Private-Group-ID = VLAN name or VLAN ID

Attribute [64] must contain the value *VLAN* (type 13). Attribute [65] must contain the value *802* (type 6). Attribute [81] specifies the *VLAN name* or *VLAN ID* assigned to the 802.1X-authenticated user.

For examples of tunnel attributes, see the [“Configuring the Switch to Use Vendor-Specific RADIUS Attributes” section on page 7-29](#).

Using 802.1X with Guest VLAN

You can configure a guest VLAN for each 802.1X port on the switch to provide limited services to clients (for example, how to download the 802.1X client). These clients might be upgrading their system for 802.1X authentication, and some hosts, such as Windows 98 systems, might not be 802.1X-capable.

When the authentication server does not receive a response to its EAPOL request/identity frame, clients that are not 802.1X-capable are put into the guest VLAN for the port, if one is configured. However, the server does not grant 802.1X-capable clients that fail authentication access to the network. Any number of hosts are allowed access when the switch port is moved to the guest VLAN. If an 802.1X-capable host joins the same port on which the guest VLAN is configured, the port is put into the unauthorized state in the user-configured access VLAN, and authentication is restarted.

Guest VLANs are supported on 802.1X ports in single-host or multiple-hosts mode.

You can configure any active VLAN except an RSPAN VLAN or a voice VLAN as an 802.1X guest VLAN. The guest VLAN feature is not supported on trunk ports; it is supported only on access ports.

For more information, see the [“Configuring a Guest VLAN” section on page 8-18](#).

Using 802.1X with Per-User ACLs

You can enable per-user access control lists (ACLs) to provide different levels of network access and service to an 802.1X-authenticated user. When the RADIUS server authenticates a user connected to an 802.1X port, it retrieves the ACL attributes based on the user identity and sends them to the switch. The switch applies the attributes to the 802.1X port for the duration of the user session. The switch removes the per-user ACL configuration when the session is over, if authentication fails, or if a link-down condition occurs. The switch does not save RADIUS-specified ACLs in the running configuration. When the port is unauthorized, the switch removes the ACL from the port.

You can configure only port ACLs on a Catalyst 2970 switch port.

RADIUS supports per-user attributes, including vendor-specific attributes. These vendor-specific attributes (VSAs) are in octet-string format and are passed to the switch during the authentication process. The VSAs used for per-user ACLs are `inACL#<n>` for the ingress direction and `outACL#<n>` for the egress direction. MAC ACLs are supported only in the ingress direction. The Catalyst 2970 switch supports VSAs only in the ingress direction. It does not support port ACLs in the egress direction on Layer 2 interfaces. For more information, see [Chapter 26, “Configuring Network Security with ACLs.”](#)

Use only the extended ACL syntax style to define the per-user configuration stored on the RADIUS server. When the definitions are passed from the RADIUS server, they are created by using the extended naming convention. However, if you use the Filter-Id attribute, it can point to a standard ACL.

You can use the Filter-Id attribute to specify an inbound or outbound ACL that is already configured on the switch. The attribute contains the ACL number followed by `.in` for ingress filtering or `.out` for egress filtering. If the RADIUS server does not allow the `.in` or `.out` syntax, the access list is applied to the outbound ACL by default. Because of limited support of Cisco IOS access lists on the switch, the Filter-Id attribute is supported only for IP ACLs numbered 1 to 199 and 1300 to 2699 (IP standard and IP extended ACLs).

Only one 802.1X-authenticated user is supported on a port. If the multiple-hosts mode is enabled on the port, the per-user ACL attribute is disabled for the associated port.

The maximum size of the per-user ACL is 4000 ASCII characters.

For examples of vendor-specific attributes, see the [“Configuring the Switch to Use Vendor-Specific RADIUS Attributes” section on page 7-29](#). For more information about configuring ACLs, see [Chapter 26, “Configuring Network Security with ACLs.”](#)

To configure per-user ACLs, you need to perform these tasks:

- Enable AAA authentication.
- Enable AAA authorization by using the **network** keyword to allow interface configuration from the RADIUS server.
- Enable 802.1X.
- Configure the user profile and VSAs on the RADIUS server.
- Configure the 802.1X port for single-host mode.

Configuring 802.1X Authentication

These sections describe how to configure 802.1X port-based authentication on your switch:

- [Default 802.1X Configuration, page 8-10](#)
- [802.1X Configuration Guidelines, page 8-11](#)

- [Upgrading from a Previous Software Release, page 8-12](#)
- [Configuring 802.1X Authentication, page 8-12](#) (required)
- [Configuring the Switch-to-RADIUS-Server Communication, page 8-13](#) (required)
- [Configuring Periodic Re-Authentication, page 8-14](#) (optional)
- [Manually Re-Authenticating a Client Connected to a Port, page 8-15](#) (optional)
- [Changing the Quiet Period, page 8-15](#) (optional)
- [Changing the Switch-to-Client Retransmission Time, page 8-16](#) (optional)
- [Setting the Switch-to-Client Frame-Retransmission Number, page 8-17](#) (optional)
- [Configuring the Host Mode, page 8-17](#) (optional)
- [Configuring a Guest VLAN, page 8-18](#) (optional)
- [Resetting the 802.1X Configuration to the Default Values, page 8-19](#) (optional)

Default 802.1X Configuration

Table 8-1 shows the default 802.1X configuration.

Table 8-1 *Default 802.1X Configuration*

Feature	Default Setting
Authentication, authorization, and accounting (AAA)	Disabled.
RADIUS server <ul style="list-style-type: none"> • IP address • UDP authentication port • Key 	<ul style="list-style-type: none"> • None specified. • 1812. • None specified.
Switch 802.1X enable state	Disabled.
Per-interface 802.1X enable state	Disabled (force-authorized). The port sends and receives normal traffic without 802.1X-based authentication of the client.
Periodic re-authentication	Disabled.
Number of seconds between re-authentication attempts	3600 seconds.
Quiet period	60 seconds (number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client).
Retransmission time	30 seconds (number of seconds that the switch should wait for a response to an EAP request/identity frame from the client before resending the request).
Maximum retransmission number	2 times (number of times that the switch will send an EAP-request/identity frame before restarting the authentication process).
Host mode	Single-host mode.

Table 8-1 Default 802.1X Configuration (continued)

Feature	Default Setting
Guest VLAN	None specified.
Client timeout period	30 seconds (when relaying a request from the authentication server to the client, the amount of time the switch waits for a response before resending the request to the client.)
Authentication server timeout period	30 seconds (when relaying a response from the client to the authentication server, the amount of time the switch waits for a reply before resending the response to the server. This setting is not configurable.)

802.1X Configuration Guidelines

These are the 802.1X authentication configuration guidelines:

- When 802.1X is enabled, ports are authenticated before any other Layer 2 feature is enabled.
- The 802.1X protocol is supported on Layer 2 static-access ports and voice VLAN ports, but it is not supported on these port types:
 - Trunk port—If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, an error message appears, and the port mode is not changed.
 - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, an error message appears, and the port mode is not changed.
 - Dynamic-access ports—If you try to enable 802.1X on a dynamic-access (VLAN Query Protocol [VQP]) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
 - EtherChannel port—Do not configure a port that is an active member of an EtherChannel as an 802.1X port. If 802.1X is enabled on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
 - Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) destination ports—You can enable 802.1X on a port that is a SPAN or RSPAN destination port. However, 802.1X is disabled until the port is removed as a SPAN or RSPAN destination port. You can enable 802.1X on a SPAN or RSPAN source port.
- You can configure any VLAN except an RSPAN VLAN or a voice VLAN as an 802.1X guest VLAN. The guest VLAN feature is not supported on trunk ports; it is supported only on access ports.
- When 802.1X is enabled on a port, you cannot configure a port VLAN that is equal to a voice VLAN.
- The 802.1X with VLAN assignment feature is not supported on trunk ports, dynamic ports, or with dynamic-access port assignment through a VMPS.

Upgrading from a Previous Software Release

In Cisco IOS Release 12.1(14)EA1, the implementation for 802.1X changed from the previous release. Some global configuration commands became interface configuration commands, and new commands were added.

If you have 802.1X configured on the switch and you upgrade to Cisco IOS Release 12.1(14)EA1 or later, the configuration file will not contain the new commands, and 802.1X will not operate. After the upgrade is complete, make sure to globally enable 802.1X by using the **dot1x system-auth-control** global configuration command. If 802.1X was running in multiple-hosts mode on an interface in the previous release, make sure to reconfigure it by using the **dot1x host-mode multi-host** interface configuration command.

Configuring 802.1X Authentication

To configure 802.1X port-based authentication, you must enable AAA and specify the authentication method list. A method list describes the sequence and authentication methods to be queried to authenticate a user.

The software uses the first method listed to authenticate users. If that method fails to respond, the software selects the next authentication method in the method list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

To allow per-user ACLs or VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

Beginning in privileged EXEC mode, follow these steps to configure 802.1X port-based authentication. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	aaa new-model	Enable AAA.
Step 3	aaa authentication dot1x {default} method1 [method2...]	<p>Create an 802.1X authentication method list.</p> <p>To create a default list that is used when a named list is <i>not</i> specified in the authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces.</p> <p>Enter at least one of these keywords:</p> <ul style="list-style-type: none"> • group radius—Use the list of all RADIUS servers for authentication. • none—Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.
Step 4	dot1x system-auth-control	Enable 802.1X authentication globally on the switch.

	Command	Purpose
Step 5	aaa authorization network {default} group radius	(Optional) Configure the switch for user RADIUS authorization for all network-related service requests, such as per-user ACLs or VLAN assignment. Note For per-user ACLs, single-host mode must be configured. This setting is the default.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the client that is to be enabled for 802.1X authentication.
Step 7	dot1x port-control auto	Enable 802.1X authentication on the interface. For feature interaction information, see the “802.1X Configuration Guidelines” section on page 8-11 .
Step 8	end	Return to privileged EXEC mode.
Step 9	show dot1x	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command. To disable 802.1X AAA authentication, use the **no aaa authentication dot1x {default | list-name}** global configuration command. To disable 802.1X AAA authorization, use the **no aaa authorization** global configuration command. To disable 802.1X authentication on the switch, use the **no dot1x system-auth-control** global configuration command.

This example shows how to enable AAA and 802.1X on Gigabit Ethernet port 0/1:

```
Switch# configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa authentication dot1x default group radius
Switch(config)# dot1x system-auth-control
Switch(config)# interface gigabitethernet0/1
Switch(config)# switchport mode access
Switch(config-if)# dot1x port-control auto
Switch(config-if)# end
```

Configuring the Switch-to-RADIUS-Server Communication

RADIUS security servers are identified by their host name or IP address, host name and specific UDP port numbers, or IP address and specific UDP port numbers. The combination of the IP address and UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the fail-over backup to the first one. The RADIUS host entries are tried in the order that they were configured.

Beginning in privileged EXEC mode, follow these steps to configure the RADIUS server parameters on the switch. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	radius-server host { <i>hostname</i> <i>ip-address</i> } auth-port <i>port-number</i> key <i>string</i>	<p>Configure the RADIUS server parameters.</p> <p>For <i>hostname</i> <i>ip-address</i>, specify the host name or IP address of the remote RADIUS server.</p> <p>For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. The default is 1812. The range is 0 to 65536.</p> <p>For key <i>string</i>, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server.</p> <p>Note Always configure the key as the last item in the radius-server host command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon.</p> <p>If you want to use multiple RADIUS servers, re-enter this command.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete the specified RADIUS server, use the **no radius-server host** {*hostname* | *ip-address*} global configuration command.

This example shows how to specify the server with IP address 172.20.39.46 as the RADIUS server, to use port 1612 as the authorization port, and to set the encryption key to *rad123*, matching the key on the RADIUS server:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1612 key rad123
```

You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using the **radius-server host** global configuration command. If you want to configure these options on a per-server basis, use the **radius-server timeout**, **radius-server retransmit**, and the **radius-server key** global configuration commands. For more information, see the [“Configuring Settings for All RADIUS Servers” section on page 7-29](#).

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch. For more information, refer to the RADIUS server documentation.

Configuring Periodic Re-Authentication

You can enable periodic 802.1X client re-authentication and specify how often it occurs. If you do not specify a time period before enabling re-authentication, the number of seconds between re-authentication attempts is 3600.

Beginning in privileged EXEC mode, follow these steps to enable periodic re-authentication of the client and to configure the number of seconds between re-authentication attempts. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	dot1x reauthentication	Enable periodic re-authentication of the client, which is disabled by default.
Step 4	dot1x timeout reauth-period <i>seconds</i>	Set the number of seconds between re-authentication attempts. The range is 1 to 65535; the default is 3600 seconds. This command affects the behavior of the switch only if periodic re-authentication is enabled.
Step 5	end	Return to privileged EXEC mode.
Step 6	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable periodic re-authentication, use the **no dot1x reauthentication** interface configuration command. To return to the default number of seconds between re-authentication attempts, use the **no dot1x timeout reauth-period** interface configuration command.

This example shows how to enable periodic re-authentication and set the number of seconds between re-authentication attempts to 4000:

```
Switch(config-if)# dot1x reauthentication
Switch(config-if)# dot1x timeout reauth-period 4000
```

Manually Re-Authenticating a Client Connected to a Port

You can manually re-authenticate the client connected to a specific port at any time by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command. This step is optional. If you want to enable or disable periodic re-authentication, see the [“Configuring Periodic Re-Authentication” section on page 8-14](#).

This example shows how to manually re-authenticate the client connected to Gigabit Ethernet port 0/1:

```
Switch# dot1x re-authenticate interface gigabitethernet0/1
```

Changing the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time and then tries again. The idle time is determined by the quiet-period value. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a smaller number than the default.

Beginning in privileged EXEC mode, follow these steps to change the quiet period. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	dot1x timeout quiet-period <i>seconds</i>	Set the number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client. The range is 1 to 65535 seconds; the default is 60.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default quiet time, use the **no dot1x timeout quiet-period** interface configuration command.

This example shows how to set the quiet time on the switch to 30 seconds:

```
Switch(config-if)# dot1x timeout quiet-period 30
```

Changing the Switch-to-Client Retransmission Time

The client responds to the EAP-request/identity frame from the switch with an EAP-response/identity frame. If the switch does not receive this response, it waits a set period of time (known as the retransmission time) and then resends the frame.



Note

You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to change the amount of time that the switch waits for client notification. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	dot1x timeout tx-period <i>seconds</i>	Set the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request. The range is 1 to 65535 seconds; the default is 30.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default retransmission time, use the **no dot1x timeout tx-period** interface configuration command.

This example shows how to set 60 as the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before resending the request:

```
Switch(config-if)# dot1x timeout tx-period 60
```

Setting the Switch-to-Client Frame-Retransmission Number

In addition to changing the switch-to-client retransmission time, you can change the number of times that the switch sends an EAP-request/identity frame (assuming no response is received) to the client before restarting the authentication process.



Note

You should change the default value of this command only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

Beginning in privileged EXEC mode, follow these steps to set the switch-to-client frame-retransmission number. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	dot1x max-req <i>count</i>	Set the number of times that the switch sends an EAP-request/identity frame to the client before restarting the authentication process. The range is 1 to 10; the default is 2.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default retransmission number, use the **no dot1x max-req** interface configuration command.

This example shows how to set 5 as the number of times that the switch sends an EAP-request/identity request before restarting the authentication process:

```
Switch(config-if)# dot1x max-req 5
```

Configuring the Host Mode

You can configure an 802.1X port for single-host or for multiple-hosts mode. In single-host mode, only one host is allowed on an 802.1X port. When the host is authenticated, the port is placed in the authorized state. When the host leaves the port, the port becomes unauthorized. Packets from hosts other than the authenticated one are dropped.

You can attach multiple hosts to a single 802.1X-enabled port as shown in [Figure 8-3 on page 8-5](#). In this mode, only one of the attached hosts must be successfully authorized for all hosts to be granted network access. If the port becomes unauthorized (re-authentication fails or an EAPOL-logoff message is received), all attached clients are denied access to the network.

With the multiple-hosts mode enabled, you can use 802.1X to authenticate the port and port security to manage network access for all MAC addresses, including that of the client.

Beginning in privileged EXEC mode, follow these steps to allow multiple hosts (clients) on an 802.1X-authorized port that has the **dot1x port-control** interface configuration command set to **auto**. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to which multiple hosts are indirectly attached.
Step 3	dot1x host-mode multi-host	Allow multiple hosts (clients) on an 802.1X-authorized port. Make sure that the dot1x port-control interface configuration command set is set to auto for the specified interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable multiple hosts on the port, use the **no dot1x host-mode multi-host** interface configuration command.

This example shows how to enable 802.1X on Gigabit Ethernet interface 0/1 and to allow multiple hosts:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# dot1x port-control auto
Switch(config-if)# dot1x host-mode multi-host
```

Configuring a Guest VLAN

When you configure a guest VLAN, clients that are not 802.1X-capable are put into the guest VLAN when the server does not receive a response to its EAPOL request/identity frame. Clients that are 802.1X-capable but fail authentication are not granted access to the network. The switch supports guest VLANs in single-host or multiple-hosts mode.

Beginning in privileged EXEC mode, follow these steps to configure a guest VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured. For the supported interface types, see the “802.1X Configuration Guidelines” section on page 8-11.
Step 3	dot1x guest-vlan <i>vlan-id</i>	Specify an active VLAN as an 802.1X guest VLAN. The range is 1 to 4094. You can configure any active VLAN except an RSPAN VLAN or a voice VLAN as an 802.1X guest VLAN.
Step 4	end	Return to privileged EXEC mode.

	Command	Purpose
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable and remove the guest VLAN, use the **no dot1x guest-vlan** interface configuration command. The port returns to the unauthorized state.

This example shows how to enable VLAN 2 as an 802.1X guest VLAN on Gigabit Ethernet interface 0/2:

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# dot1x guest-vlan 2
```

Resetting the 802.1X Configuration to the Default Values

Beginning in privileged EXEC mode, follow these steps to reset the 802.1X configuration to the default values. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	dot1x default	Reset the configurable 802.1X parameters to the default values.
Step 4	end	Return to privileged EXEC mode.
Step 5	show dot1x interface <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Displaying 802.1X Statistics and Status

To display 802.1X statistics for all interfaces, use the **show dot1x all statistics** privileged EXEC command. To display 802.1X statistics for a specific interface, use the **show dot1x statistics interface** *interface-id* privileged EXEC command.

To display the 802.1X administrative and operational status for the switch, use the **show dot1x all** privileged EXEC command. To display the 802.1X administrative and operational status for a specific interface, use the **show dot1x interface** *interface-id* privileged EXEC command.

For detailed information about the fields in these displays, refer to the command reference for this release.



Configuring Interface Characteristics

This chapter defines the types of interfaces on the Catalyst 2970 switch and describes how to configure them.

The chapter has these sections:

- [Understanding Interface Types, page 9-1](#)
- [Using Interface Configuration Mode, page 9-5](#)
- [Configuring Ethernet Interfaces, page 9-9](#)
- [Configuring the System MTU, page 9-14](#)
- [Monitoring and Maintaining the Interfaces, page 9-16](#)



Note

For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and the online *Cisco IOS Interface Command Reference for Release 12.1*.

Understanding Interface Types

This section describes the different types of interfaces supported by the switch with references to chapters that contain more detailed information about configuring these interface types. The rest of the chapter describes configuration procedures for physical interface characteristics.

These sections are included:

- [Port-Based VLANs, page 9-2](#)
- [Switch Ports, page 9-2](#)
- [EtherChannel Port Groups, page 9-4](#)
- [Connecting Interfaces, page 9-4](#)

Port-Based VLANs

A VLAN is a switched network that is logically segmented by function, team, or application, without regard to the physical location of the users. For more information about VLANs, see [Chapter 11, “Configuring VLANs.”](#) Packets received on a port are forwarded only to ports that belong to the same VLAN as the receiving port. Network devices in different VLANs cannot communicate with one another without a Layer 3 device to route traffic between the VLANs.

VLAN partitions provide hard firewalls for traffic in the VLAN, and each VLAN has its own MAC address table. A VLAN comes into existence when a local port is configured to be associated with the VLAN, when the VLAN Trunking Protocol (VTP) learns of its existence from a neighbor on a trunk, or when a user creates a VLAN.

To configure normal-range VLANs (VLAN IDs 1 to 1005), use the **vlan *vlan-id*** global configuration command to enter config-vlan mode or the **vlan database** privileged EXEC command to enter VLAN database configuration mode. The VLAN configurations for VLAN IDs 1 to 1005 are saved in the VLAN database. To configure extended-range VLANs (VLAN IDs 1006 to 4094), you must use config-vlan mode with VTP mode set to transparent. Extended-range VLANs are not added to the VLAN database. When VTP mode is transparent, the VTP and VLAN configuration is saved in the switch running configuration, and you can save it in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command.

Add ports to a VLAN by using the **switchport** interface configuration commands:

- Identify the interface.
- For a trunk port, set trunk characteristics, and if desired, define the VLANs to which it can belong.
- For an access port, set and define the VLAN to which it belongs.

Switch Ports

Switch ports are Layer 2-only interfaces associated with a physical port. Switch ports belong to one or more VLANs. A switch port can be an access port or a trunk port. You can configure a port as an access port or trunk port or let the Dynamic Trunking Protocol (DTP) operate on a per-port basis to determine switchport mode by negotiating with the port on the other end of the link. Switch ports are used for managing the physical interface and associated Layer 2 protocols.

Configure switch ports by using the **switchport** interface configuration commands. For detailed information about configuring access port and trunk port characteristics, see [Chapter 11, “Configuring VLANs.”](#)

Access Ports

An access port belongs to and carries the traffic of only one VLAN (unless it is configured as a voice VLAN port). Traffic is received and sent in native formats with no VLAN tagging. Traffic arriving on an access port is assumed to belong to the VLAN assigned to the port. If an access port receives a tagged packet (Inter-Switch Link [ISL] or 802.1Q tagged), the packet is dropped, and the source address is not learned.

Two types of access ports are supported:

- Static access ports are manually assigned to a VLAN.
- VLAN membership of dynamic access ports is learned through incoming packets. By default, a dynamic access port is a member of no VLAN, and forwarding to and from the port is enabled only when the VLAN membership of the port is discovered. Dynamic access ports on the switch are assigned to a VLAN by a VLAN Membership Policy Server (VMPS). The VMPS can be a Catalyst 6000 series switch; the Catalyst 2970 switch cannot be a VMPS server.

You can also configure an access port with an attached Cisco IP Phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see [Chapter 13, “Configuring Voice VLAN.”](#)

Trunk Ports

A trunk port carries the traffic of multiple VLANs and by default is a member of all VLANs in the VLAN database. Two types of trunk ports are supported:

- In an ISL trunk port, all received packets are expected to be encapsulated with an ISL header, and all transmitted packets are sent with an ISL header. Native (non-tagged) frames received from an ISL trunk port are dropped.
- An IEEE 802.1Q trunk port supports simultaneous tagged and untagged traffic. An 802.1Q trunk port is assigned a default Port VLAN ID (PVID), and all untagged traffic travels on the port default PVID. All untagged traffic and tagged traffic with a NULL VLAN ID are assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.

Although by default, a trunk port is a member of every VLAN known to the VTP, you can limit VLAN membership by configuring an allowed list of VLANs for each trunk port. The list of allowed VLANs does not affect any other port but the associated trunk port. By default, all possible VLANs (VLAN ID 1 to 4094) are in the allowed list. A trunk port can only become a member of a VLAN if VTP knows of the VLAN and the VLAN is in the enabled state. If VTP learns of a new, enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of that VLAN and traffic is forwarded to and from the trunk port for that VLAN. If VTP learns of a new, enabled VLAN that is not in the allowed list for a trunk port, the port does not become a member of the VLAN, and no traffic for the VLAN is forwarded to or from the port.

For more information about trunk ports, see [Chapter 11, “Configuring VLANs.”](#)

EtherChannel Port Groups

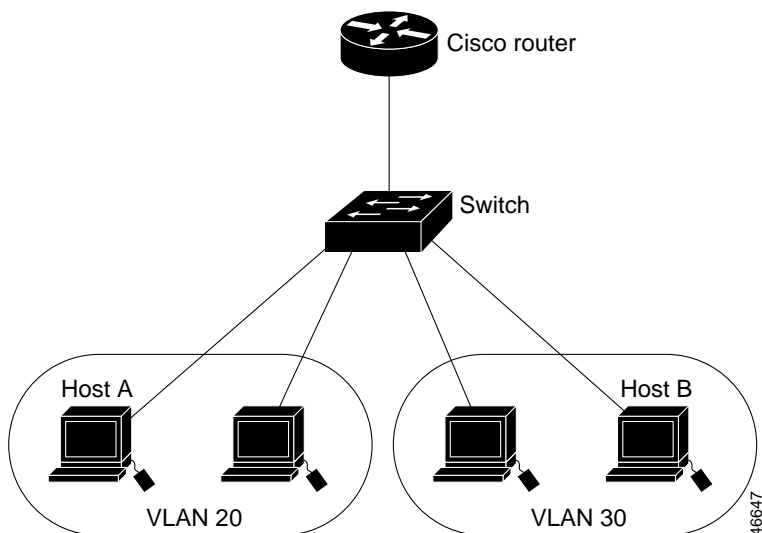
EtherChannel port groups provide the ability to treat multiple switch ports as one switch port. These port groups act as a single logical port for high-bandwidth connections between switches or between switches and servers. An EtherChannel balances the traffic load across the links in the channel. If a link within the EtherChannel fails, traffic previously carried over the failed link changes to the remaining links. You can group multiple trunk ports into one logical trunk port or multiple access ports into one logical access port. Most protocols operate over either single ports or aggregated switch ports and do not recognize the physical ports within the port group. Exceptions are the DTP, the Cisco Discovery Protocol (CDP), and the Port Aggregation Protocol (PAgP), which operate only on physical ports.

When you configure an EtherChannel, you create a port-channel logical interface and assign an interface to the EtherChannel. Use the **channel-group** interface configuration command to dynamically create the port-channel logical interface. This command binds the physical and logical ports together. For more information, see [Chapter 28, “Configuring EtherChannels.”](#)

Connecting Interfaces

Devices within a single VLAN can communicate directly through any switch. Ports in different VLANs cannot exchange data without going through a routing device. In the configuration shown in [Figure 9-1](#), when Host A in VLAN 20 sends data to Host B in VLAN 30, it must go from Host A to the switch, to the router, back to the switch, and then to Host B.

Figure 9-1 Connecting VLANs with Layer 2 Switches



Using Interface Configuration Mode

The switch supports these interface types:

- Physical ports—including switch ports and routed ports
- VLANs—switch virtual interfaces
- Port-channels—EtherChannel of interfaces

You can also configure a range of interfaces (see the [“Configuring a Range of Interfaces”](#) section on page 9-6).

To configure a physical interface (port), enter interface configuration mode, and specify the interface type, module number, and switch port number.

- Type—Gigabit Ethernet (gigabitethernet or gi) for 10/100/1000 Mbps Ethernet ports or small form-factor pluggable (SFP) Gigabit Ethernet interfaces.
- Module number—The module or slot number on the switch (always 0 on the Catalyst 2970 switch).
- Port number—The interface number on the switch. The port numbers always begin at 1, starting at the left when facing the front of the switch, for example, gigabitethernet0/1, gigabitethernet0/2. On a switch with SFP modules, the SFP module ports are numbered consecutively following the 10/100/1000 interfaces. The SFP module ports are gigabitethernet 0/25 through gigabitethernet 0/28.

You can identify physical interfaces by physically checking the interface location on the switch. You can also use the Cisco IOS **show** privileged EXEC commands to display information about a specific interface or all the interfaces on the switch. The remainder of this chapter primarily provides physical interface configuration procedures.

Procedures for Configuring Interfaces

These general instructions apply to all interface configuration processes.

- Step 1** Enter the **configure terminal** command at the privileged EXEC prompt:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#
```

- Step 2** Enter the **interface** global configuration command. Identify the interface type and the number of the connector. In this example, Gigabit Ethernet interface 0/1 is selected:

```
Switch(config)# interface GigabitEthernet0/1
Switch(config-if)#
```



Note You do not need to add a space between the interface type and interface number. For example, in the preceding line, you can specify either **gigabitethernet 0/1**, **gigabitethernet0/1**, **gi 0/1**, or **gi0/1**.

- Step 3** Follow each **interface** command with the interface configuration commands that the interface requires. The commands that you enter define the protocols and applications that will run on the interface. The commands are collected and applied to the interface when you enter another interface command or enter **end** to return to privileged EXEC mode.

You can also configure a range of interfaces by using the **interface range** or **interface range macro** global configuration commands. Interfaces configured in a range must be the same type and must be configured with the same feature options.

Step 4 After you configure an interface, verify its status by using the **show** privileged EXEC commands listed in the [“Monitoring and Maintaining the Interfaces”](#) section on page 9-16.

Enter the **show interfaces** privileged EXEC command to see a list of all interfaces on or configured for the switch. A report is provided for each interface that the device supports or for the specified interface.

Configuring a Range of Interfaces

You can use the **interface range** global configuration command to configure multiple interfaces with the same configuration parameters. When you enter the interface range configuration mode, all command parameters that you enter are attributed to all interfaces within that range until you exit this mode.

Beginning in privileged EXEC mode, follow these steps to configure a range of interfaces with the same parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface range { <i>port-range</i> macro <i>macro_name</i> }	Enter interface range configuration mode by entering the range of interfaces (VLANs or physical ports) to be configured. <ul style="list-style-type: none"> You can use the interface range command to configure up to five port ranges or a previously defined macro. The macro variable is explained in the “Configuring and Using Interface Range Macros” section on page 9-7. In a comma-separated <i>port-range</i>, you must enter the interface type for each entry and enter spaces before and after the comma. In a hyphen-separated <i>port-range</i>, you do not need to re-enter the interface type, but you must enter a space before the hyphen.
Step 3		You can now use the normal configuration commands to apply the configuration parameters to all interfaces in the range.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces [<i>interface-id</i>]	Verify the configuration of the interfaces in the range.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When using the **interface range** global configuration command, note these guidelines:

- Valid entries for *port-range*:
 - vlan** *vlan-ID* - *vlan-ID*, where VLAN ID is from 1 to 4094
 - gigabitethernet** *module*/{*first port*} - {*last port*}, where the module is 0
 - port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 12



Note When you use the **interface range** command with port channels, the first and last port channel number must be active port channels.

- You must add a space between the first interface number and the hyphen when using the **interface range** command. For example, the command **interface range gigabitethernet 0/1 - 4** is a valid range; the command **interface range gigabitethernet 0/1-4** is not a valid range.
- The **interface range** command only works with VLAN interfaces that have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used with the **interface range** command.
- All interfaces defined as in a range must be the same type, all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs, but you can enter multiple ranges in a command.

This example shows how to use the **interface range** global configuration command to set the speed on 10/100/1000 interfaces 0/1 to 0/4 to 100 Mbps:

```
Switch# configure terminal
Switch(config)# interface range GigabitEthernet0/1 - 4
Switch(config-if-range)# speed 100
```

If you enter multiple configuration commands while you are in interface range mode, each command is executed as it is entered. The commands are not batched together and executed after you exit interface range mode. If you exit interface range configuration mode while the commands are being executed, some commands might not be executed on all interfaces in the range. Wait until the command prompt reappears before exiting interface range configuration mode.

Configuring and Using Interface Range Macros

You can create an interface range macro to automatically select a range of interfaces for configuration. Before you can use the **macro** keyword in the **interface range macro** global configuration command string, you must use the **define interface-range** global configuration command to define the macro.

Beginning in privileged EXEC mode, follow these steps to define an interface range macro:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	define interface-range <i>macro_name</i> <i>interface-range</i>	Define the interface-range macro, and save it in NVRAM. <ul style="list-style-type: none"> • The <i>macro_name</i> is a 32-character maximum character string. • A macro can contain up to five comma-separated interface ranges. • Each <i>interface-range</i> must consist of the same port type.
Step 3	interface range macro <i>macro_name</i>	Select the interface range to be configured using the values saved in the interface-range macro called <i>macro_name</i> . You can now use the normal configuration commands to apply the configuration to all interfaces in the defined macro.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config include define	Show the defined interface range macro configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no define interface-range** *macro_name* global configuration command to delete a macro.

When using the **define interface-range** global configuration command, note these guidelines:

- Valid entries for *interface-range*:
 - **vlan** *vlan-ID* - *vlan-ID*, where VLAN ID is from 1 to 4094
 - **gigabitethernet** *module*/{*first port*} - {*last port*}, where the module is **0**
 - **port-channel** *port-channel-number* - *port-channel-number*, where *port-channel-number* is from 1 to 12.



Note

When you use the interface ranges with port channels, the first and last port channel number must be active port channels.

- You must add a space between the first interface number and the hyphen when entering an *interface-range*. For example, **gigabitethernet 0/1 - 4** is a valid range; **gigabitethernet 0/1-4** is not a valid range.
- The VLAN interfaces must have been configured with the **interface vlan** command. The **show running-config** privileged EXEC command displays the configured VLAN interfaces. VLAN interfaces not displayed by the **show running-config** command cannot be used as *interface-ranges*.
- All interfaces defined as in a range must be the same type (all Gigabit Ethernet ports, all EtherChannel ports, or all VLANs), but you can combine multiple interface types in a macro.

This example shows how to define an interface-range macro named *enet_list* to select Gigabit Ethernet ports 0/1 to 0/4 and to verify the macro configuration:

```
Switch# configure terminal
Switch(config)# define interface-range enet_list GigabitEthernet0/1 - 4
Switch(config)# end
Switch# show running-config | include define
define interface-range enet_list GigabitEthernet0/1 - 4
```

This example shows how to create a multiple-interface macro named *macro1*:

```
Switch# configure terminal
Switch(config)# define interface-range macro1 GigabitEthernet0/1 - 2,
GigabitEthernet0/5 - 7
Switch(config)# end
```

This example shows how to enter interface range configuration mode for the interface-range macro *enet_list*:

```
Switch# configure terminal
Switch(config)# interface range macro enet_list
Switch(config-if-range)#
```

This example shows how to delete the interface-range macro *enet_list* and to verify that it was deleted.

```
Switch# configure terminal
Switch(config)# no define interface-range enet_list
Switch(config)# end
Switch# show run | include define
Switch#
```


Configuring Ethernet Interfaces

These sections describe the default interface configuration and the optional features that you can configure on most physical interfaces:

- [Default Ethernet Interface Configuration, page 9-9](#)
- [Configuring Interface Speed and Duplex Mode, page 9-10](#)
- [Configuring IEEE 802.3X Flow Control, page 9-12](#)
- [Configuring Auto-MDIX on an Interface, page 9-13](#)
- [Adding a Description for an Interface, page 9-14](#)

Default Ethernet Interface Configuration

[Table 9-1](#) shows the Ethernet interface default configuration. For more details on the VLAN parameters listed in the table, see [Chapter 11, “Configuring VLANs.”](#) For details on controlling traffic to the port, see [Chapter 19, “Configuring Port-Based Traffic Control.”](#)

Table 9-1 *Default Layer 2 Ethernet Interface Configuration*

Feature	Default Setting
Allowed VLAN range	VLANs 1 – 4094.
Default VLAN (for access ports)	VLAN 1.
Native VLAN (for 802.1Q trunks)	VLAN 1.
VLAN trunking	Switchport mode dynamic auto (supports DTP) (Layer 2 interfaces only).
Port enable state	All ports are enabled.
Port description	None defined.
Speed	Autonegotiate.
Duplex mode	Autonegotiate.
Flow control	Flow control is set to receive: off . It is always off for sent packets.
EtherChannel (PAgP)	Disabled on all Ethernet ports. See Chapter 28, “Configuring EtherChannels.”
Port blocking (unknown multicast and unknown unicast traffic)	Disabled (not blocked) (Layer 2 interfaces only). See the “Configuring Port Blocking” section on page 19-6.
Broadcast, multicast, and unicast storm control	Disabled. See the “Default Storm Control Configuration” section on page 19-3.
Protected port	Disabled. See the “Configuring Protected Ports” section on page 19-5.
Port security	Disabled. See the “Default Port Security Configuration” section on page 19-9. L2

Table 9-1 Default Layer 2 Ethernet Interface Configuration (continued)

Feature	Default Setting
Port Fast	Disabled.
Auto-MDIX	Disabled. Note The switch might not support a pre-standard power device—such as Cisco IP phones and access points that do not fully support IEEE 802.3AF—if that power device is connected to the switch through a crossover cable. This is regardless of whether Auto-MIDX is enabled on the switch port.

Configuring Interface Speed and Duplex Mode

Ethernet interfaces on the switch operate at 10, 100, or 1000 Mbps and in either full- or half-duplex mode. In full-duplex mode, two stations can send and receive traffic at the same time. Normally, 10-Mbps ports operate in half-duplex mode, which means that stations can either receive or send traffic.

Switch models include Gigabit Ethernet (10/100/1000-Mbps) ports and small form-factor pluggable (SFP) module slots supporting Gigabit SFP modules.

- You can configure interface speed on Gigabit Ethernet (10/100/1000-Mbps) interfaces. You can configure Gigabit Ethernet interfaces to full-duplex mode or to autonegotiate; you cannot configure half-duplex mode on Gigabit Ethernet ports.
- You cannot configure speed or duplex mode on most SFP ports, but you can configure speed to not negotiate (**nonegotiate**) if connected to a device that does not support autonegotiation. However, when a 1000BASE-T SFP module is in the SFP module port, you can configure speed as 10, 100, or 1000 Mbps, and you can configure duplex mode to auto or full.

These sections describe how to configure the interface speed and duplex mode:

- [Configuration Guidelines, page 9-10](#)
- [Setting the Interface Speed and Duplex Parameters, page 9-11](#)

Configuration Guidelines

When configuring an interface speed and duplex mode, note these guidelines:

- If both ends of the line support autonegotiation, we highly recommend the default setting of **auto** negotiation.
- If one interface supports autonegotiation and the other end does not, configure duplex and speed on both interfaces; do not use the **auto** setting on the supported side.
- For 10/100/1000 Mbps ports, if both the speed and duplex mode are set to specific values, autonegotiation is disabled.
- You cannot configure duplex mode on SFP module ports; they operate in full-duplex mode. However, when a 1000BASE-T SFP module is inserted in an SFP module port, you can configure the duplex mode to **full** or **auto** and half-duplex mode is supported with the auto configuration.
- You cannot configure speed on SFP module ports, except to **nonegotiate**. However, when a 1000BASE-T SFP module is in the SFP module port, the speed can be configured to **10**, **100**, **1000**, or **auto**, but not **nonegotiate**.

- When STP is enabled and a port is reconfigured, the switch can take up to 30 seconds to check for loops. The port LED is amber while STP reconfigures.

**Caution**

Changing the interface speed and duplex mode configuration might shut down and re-enable the interface during the reconfiguration.

Setting the Interface Speed and Duplex Parameters

Beginning in privileged EXEC mode, follow these steps to set the speed and duplex mode for a physical interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface identification.
Step 3	speed { 10 100 1000 auto nonegotiate }	<p>Enter the appropriate speed parameter for the interface:</p> <ul style="list-style-type: none"> • Enter 10, 100, or 1000 to set a specific speed for the interface. Enter auto to enable the interface to autonegotiate speed with the device connected to the interface. • The nonegotiate keyword is available only for SFP module ports. SFP module ports operate only at 1000 Mbps but can be configured to not negotiate if connected to a device that does not support autonegotiation. <p>Note When a 1000BASE-T SFP module is in the SFP module port, the speed can be configured to 10, 100, 1000, or auto, but not nonegotiate.</p>
Step 4	duplex { auto full }	<p>Enter the duplex parameter for the interface.</p> <p>You cannot configure Gigabit Ethernet interfaces to operate in half-duplex mode.</p> <p>This command is not available on SFP module ports unless a 1000BASE-T SFP module is inserted, when the mode can then be configured to auto or full.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i>	Display the interface speed and duplex mode configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no speed** and **no duplex** interface configuration commands to return the interface to the default speed and duplex settings (autonegotiate). To return all interface settings to the defaults, use the **default interface** *interface-id* interface configuration command.

This example shows how to set the interface speed to 100 Mbps on Gigabit Ethernet interface 0/3:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# speed 100
```

Configuring IEEE 802.3X Flow Control

Flow control enables connected Ethernet ports to control traffic rates during congestion by allowing congested nodes to pause link operation at the other end. If one port experiences congestion and cannot receive any more traffic, it notifies the other port to stop sending until the condition clears by sending a pause frame. Upon receipt of a pause frame, the sending device stops sending any data packets, which prevents any loss of data packets during the congestion period.



Note

Catalyst 2970 ports are capable of receiving, but not sending, pause frames.

You use the **flowcontrol** interface configuration command to set the interface's ability to **receive** pause frames to **on**, **off**, or **desired**. The default state is **off**.

When set to **desired**, an interface can operate with an attached device that is required to send flow-control packets or with an attached device that is not required to but can send flow-control packets.

These rules apply to flow control settings on the device:

- **receive on** (or **desired**): The port cannot send pause frames but can operate with an attached device that is required to or can send pause frames; the port can receive pause frames.
- **receive off**: Flow control does not operate in either direction. In case of congestion, no indication is given to the link partner, and no pause frames are sent or received by either device.



Note

For details on the command settings and the resulting flow control resolution on local and remote ports, refer to the **flowcontrol** interface configuration command in the command reference for this release.

Beginning in privileged EXEC mode, follow these steps to configure flow control on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the physical interface to be configured.
Step 3	flowcontrol { receive } { on off desired }	Configure the flow control mode for the port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i>	Verify the interface flow control settings.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable flow control, use the **flowcontrol receive off** interface configuration command.

This example shows how to turn on flow control on Gigabit Ethernet interface 0/1:

```
Switch# configure terminal
Switch(config)# interface GigabitEthernet0/1
Switch(config-if)# flowcontrol receive on
Switch(config-if)# end
```

Configuring Auto-MDIX on an Interface

When automatic medium-dependent interface crossover (Auto-MDIX) is enabled on an interface, the interface automatically detects the required cable connection type (straight through or crossover) and configures the connection appropriately. When connecting switches without the Auto-MDIX feature, you must use straight-through cables to connect to devices such as servers, workstations, or routers and crossover cables to connect to other switches or repeaters. With Auto-MDIX enabled, you can use either type of cable to connect to other devices, and the interface automatically corrects for any incorrect cabling. For more information about cabling requirements, refer to the hardware installation guide.

Auto-MDIX is disabled by default. When you enable Auto-MDIX, you must also set the speed and duplex on the interface to **auto** in order for the feature to operate correctly. Auto-MDIX is supported on all 10/100/1000 Mbps interfaces and on 10/100/1000 BASE-T/TX SFP interfaces. It is not supported on 1000 BASE-SX or -LX SFP interfaces.

Table 9-2 shows the link states that results from Auto-MDIX settings and correct and incorrect cabling.

Table 9-2 Link Conditions and Auto-MDIX Settings

Local Side Auto-MDIX	Remote Side Auto-MDIX	With Correct Cabling	With Incorrect Cabling
On	On	Link up	Link up
On	Off	Link up	Link up
Off	On	Link up	Link up
Off	Off	Link up	Link down

Beginning in privileged EXEC mode, follow these steps to configure Auto-MDIX on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the physical interface to be configured.
Step 3	speed auto	Configure the interface to autonegotiate speed with the connected device.
Step 4	duplex auto	Configure the interface to autonegotiate duplex mode with the connected device.
Step 5	mdix auto	Enable Auto-MDIX on the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show controllers ethernet-controller <i>interface-id</i> phy	Verify the operational state of the Auto-MDIX feature on the interface.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable Auto-MDIX, use the **no mdix auto** interface configuration command.

This example shows how to enable Auto-MDIX on Gigabit Ethernet interface 0/1:

```
Switch# configure terminal
Switch(config)# interface GigabitEthernet0/1
Switch(config-if)# speed auto
Switch(config-if)# duplex auto
Switch(config-if)# mdix auto
Switch(config-if)# end
```

Adding a Description for an Interface

You can add a description about an interface to help you remember its function. The description appears in the output of these privileged EXEC commands: **show configuration**, **show running-config**, and **show interfaces**.

Beginning in privileged EXEC mode, follow these steps to add a description for an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface for which you are adding a description.
Step 3	description <i>string</i>	Add a description (up to 240 characters) for an interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> description or show running-config	Verify your entry.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no description** interface configuration command to delete the description.

This example shows how to add a description on Gigabit Ethernet interface 0/3 and to verify the description:

```
Switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface GigabitEthernet0/3
Switch(config-if)# description Connects to Marketing
Switch(config-if)# end
Switch# show interfaces GigabitEthernet0/3 description
Interface Status      Protocol Description
Gi0/3      admin down      down    Connects to Marketing
```

Configuring the System MTU

The default maximum transmission unit (MTU) size for frames received and transmitted on all interfaces on the switch is 1500 bytes. You can increase the MTU size for all interfaces operating at 10 or 100 Mbps by using the **system mtu** global configuration command. You can increase the MTU size to support jumbo frames on all Gigabit Ethernet interfaces by using the **system mtu jumbo** global configuration command. Gigabit Ethernet ports are not affected by the **system mtu** command; 10/100 ports are not affected by the **system jumbo mtu** command.

You cannot set the MTU size for an individual interface; you set it for all 10/100 or all Gigabit Ethernet interfaces on the switch. When you change the MTU size, you must reset the switch before the new configuration takes effect.

The size of frames that can be received by the switch CPU is limited to 1500 bytes, no matter what value was entered with the **system mtu** or **system mtu jumbo** commands. Although frames that are forwarded typically are not received by the CPU, in some cases packets are sent to the CPU, such as traffic sent to control traffic, SNMP, or Telnet.

**Note**

If Gigabit Ethernet interfaces are configured to accept frames greater than the 10/100 interfaces, jumbo frames ingressing on a Gigabit Ethernet interface and egressing on a 10/100 interface are dropped.

Beginning in privileged EXEC mode, follow these steps to change MTU size for all 10/100 or Gigabit Ethernet interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	system mtu <i>bytes</i>	(Optional) Change the MTU size for all interfaces on the switch that are operating at 10 or 100 Mbps. The range is from 1500 to 1546 bytes; the default is 1500 bytes.
Step 3	system mtu jumbo <i>bytes</i>	(Optional) Change the MTU size for all Gigabit Ethernet interfaces on the switch. The range is from 1500 to 9000 bytes; the default is 1500 bytes.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	Save your entries in the configuration file.
Step 6	reload	Reload the operating system.

If you enter a value that is outside the allowed range for the specific type of interface, the value is not accepted.

Once the switch reloads, you can verify your settings by entering the **show system mtu** privileged EXEC command.

This example shows how to set the maximum packet size for a Gigabit Ethernet port to 1800 bytes:

```
Switch(config)# system jumbo mtu 1800
Switch(config)# exit
Switch# reload
```

This example shows the response when you try to set Gigabit Ethernet interfaces to an out-of-range number:

```
Switch(config)# system mtu jumbo 2500
                        ^
% Invalid input detected at '^' marker.
```

Monitoring and Maintaining the Interfaces

You can perform the tasks in these sections to monitor and maintain interfaces:

- [Monitoring Interface Status, page 9-16](#)
- [Clearing and Resetting Interfaces and Counters, page 9-17](#)
- [Shutting Down and Restarting the Interface, page 9-17](#)

Monitoring Interface Status

Commands entered at the privileged EXEC prompt display information about the interface, including the versions of the software and the hardware, the configuration, and statistics about the interfaces. [Table 9-3](#) lists some of these interface monitoring commands. (You can display the full list of **show** commands by using the **show ?** command at the privileged EXEC prompt.) These commands are fully described in the *Cisco IOS Interface Command Reference for Release 12.1*.

Table 9-3 Show Commands for Interfaces

Command	Purpose
show interfaces [<i>interface-id</i>]	Display the status and configuration of all interfaces or a specific interface.
show interfaces <i>interface-id</i> status [err-disabled]	Display interface status or a list of interfaces in an error-disabled state.
show interfaces [<i>interface-id</i>] switchport	Display administrative and operational status of switching ports.
show interfaces [<i>interface-id</i>] description	Display the description configured on an interface or all interfaces and the interface status.
show ip interface [<i>interface-id</i>]	Display the usability status of all interfaces configured for IP routing or the specified interface.
show running-config interface [<i>interface-id</i>]	Display the running configuration in RAM for the interface.
show version	Display the hardware configuration, software version, the names and sources of configuration files, and the boot images.
show controllers ethernet-controller <i>interface-id</i> phy	Verify the operational state of the Auto-MDIX feature on the interface.

Clearing and Resetting Interfaces and Counters

Table 9-4 lists the privileged EXEC mode **clear** commands that you can use to clear counters and reset interfaces.

Table 9-4 Clear Commands for Interfaces

Command	Purpose
clear counters [<i>interface-id</i>]	Clear interface counters.
clear interface <i>interface-id</i>	Reset the hardware logic on an interface.
clear line [<i>number</i> console 0 vty number]	Reset the hardware logic on an asynchronous serial line.

To clear the interface counters shown by the **show interfaces** privileged EXEC command, use the **clear counters** privileged EXEC command. The **clear counters** command clears all current interface counters from the interface unless optional arguments are specified to clear only a specific interface type from a specific interface number.



Note

The **clear counters** privileged EXEC command does not clear counters retrieved by using Simple Network Management Protocol (SNMP), but only those seen with the **show interface** privileged EXEC command.

Shutting Down and Restarting the Interface

Shutting down an interface disables all functions on the specified interface and marks the interface as unavailable on all monitoring command displays. This information is communicated to other network servers through all dynamic routing protocols. The interface is not mentioned in any routing updates.

Beginning in privileged EXEC mode, follow these steps to shut down an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface { vlan <i>vlan-id</i> } { { fastethernet gigabitethernet } <i>interface-id</i> } { port-channel <i>port-channel-number</i> }	Select the interface to be configured.
Step 3	shutdown	Shut down an interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entry.

Use the **no shutdown** interface configuration command to restart the interface.

To verify that an interface is disabled, enter the **show interfaces** privileged EXEC command. A disabled interface is shown as *administratively down* in the **show interface** command display.



Configuring SmartPort Macros

This chapter describes how to configure and apply SmartPort macros on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding SmartPort Macros, page 10-1](#)
- [Configuring Smart-Port Macros, page 10-2](#)
- [Displaying SmartPort Macros, page 10-4](#)

Understanding SmartPort Macros

SmartPort macros provide a convenient way to save and share common configurations. You can use SmartPort macros to enable features and settings based on the location of a switch in the network and for mass configuration deployments across the network.

Each SmartPort macro is a set of CLI commands that you define. SmartPort macros do not contain new CLI commands; they are simply a group of existing CLI commands.

When you apply a SmartPort macro on an interface, the CLI commands contained within the macro are configured on the interface. When the macro is applied to an interface, the existing interface configurations are not lost. The new commands are added to interface and are saved in the running configuration file.

Configuring Smart-Port Macros

You can create a new SmartPort macro or use an existing macro as a template to create a new macro that is specific to your application. After you create the macro, you can apply it to an interface or range of interfaces.

This section includes information about:

- [Default SmartPort Macro Configuration, page 10-2](#)
- [SmartPort Macro Configuration Guidelines, page 10-2](#)
- [Creating and Applying SmartPort Macros, page 10-3](#)

Default SmartPort Macro Configuration

There are no default SmartPort macros configured on the switch.

SmartPort Macro Configuration Guidelines

Follow these guidelines when configuring macros on your switch:

- Do not use **exit** or **end** commands when creating a macro. This could cause commands that follow **exit** or **end** to execute in a different command mode.
- When creating a macro, all CLI commands should be interface configuration mode commands.
- Some CLI commands are specific to certain interface types. The macro will fail the syntax check or the configuration check, and the switch will return an error message if it is applied to an interface that does not accept the configuration.
- When a macro is applied to an interface, all existing configuration on the interface is retained. This is helpful when applying an incremental configuration to an interface.
- If you modify a macro definition by adding or deleting commands, the changes are not reflected on the interface where the original macro was applied. You need to reapply the updated macro on the interface to apply the new or changed commands.
- You can use the **macro trace macro-name** interface configuration command to show what macros are running on an interface or to debug the macro to determine any syntax or configuration errors.
- If a command fails when you apply a macro, either due to a syntax error or to a configuration error, the macro continues to apply the remaining commands to the interface.
- Applying a macro to an interface range is the same as applying a macro to a single interface. When you use an interface range, the macro is applied sequentially to each individual interface within the range. If a macro command fails on one interface, it is still applied to the remaining interfaces.

Creating and Applying SmartPort Macros

Beginning in privileged EXEC mode, follow these steps to create and apply a SmartPort macro:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	macro name <i>macro-name</i>	Create a macro definition, and enter a macro name. A macro definition can contain up to 3000 characters. Enter the macro commands with one command per line. Use the @ character to end the macro. Use the # character at the beginning of a line to enter comment text within the macro. We recommend that you do not use the exit or end commands in a macro. This could cause any commands following exit or end to execute in a different command mode. For best results, all commands in a macro should be interface configuration mode commands.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to apply the macro.
Step 4	macro {apply trace} <i>macro-name</i>	Apply each individual command defined in the macro to the interface by entering macro apply macro-name . Specify macro trace macro-name to apply and print each command before it is applied to the interface.
Step 5	macro description <i>text</i>	(Optional) Enter a description about the macro that is applied to the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show parser macro	Verify that the macro was created.
Step 8	show running-config interface <i>interface-id</i>	Verify that the macro is applied to an interface.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **no** form of the **macro name** global configuration command only deletes the macro definition. It does not affect the configuration of those interfaces on which the macro is already applied. You can delete a macro-applied configuration on an interface by entering the **default interface interface-id** interface configuration command. Alternatively, you can create an *anti-macro* for an existing macro that contains the **no** form of all the corresponding commands in the original macro. Then apply the anti-macro to the interface.

This example shows how to define the **desktop-config** macro for an access switch interface, apply the macro to Gigabit Ethernet interface 0/4, add a description to the interface, and verify the configuration.

```
Switch(config)# macro name desktop-config
# Put the switch in access mode
switchport mode access
# Allow port to move to forwarding state quickly
spanning-tree portfast
# BPDUs should not be sent into the network
spanning-tree bpduguard enable
# Restrict the port to one address -- that of desktop
switchport port-security maximum 1
# Put all data traffic in vlan 1
switchport access vlan 1
@

Switch(config)# interface gigabitethernet0/4
Switch(config-if)# macro apply desktop-config
Switch(config-if)# macro description desktop-config
Switch(config-if)# end
Switch# show parser macro name desktop-config
Macro name : desktop-config
Macro type : customizable

macro description desktop-config
# Put the switch in access mode
switchport mode access
# Allow port to move to forwarding state quickly
spanning-tree portfast
# BPDUs should not be sent into the network
spanning-tree bpduguard enable
# Restrict the port to one address -- that of desktop
switchport port-security maximum 1
# Put all data traffic in vlan 1
switchport access vlan 1

Switch# show parser macro description
Interface      Macro Description
-----
Gi0/9          desktop-config
-----
```

Displaying SmartPort Macros

To display the SmartPort macros, use one or more of the privileged EXEC commands in [Table 10-1](#).

Table 10-1 Commands for Displaying SmartPort Macros

Command	Purpose
show parser macro	Displays all configured macros.
show parser macro name <i>macro-name</i>	Displays a specific macro.
show parser macro brief	Displays the configured macro names.
show parser macro description [interface <i>interface-id</i>]	Displays the macro description for all interfaces or for a specified interface.



Configuring VLANs

This chapter describes how to configure normal-range VLANs (VLAN IDs 1 to 1005) and extended-range VLANs (VLAN IDs 1006 to 4094) on the Catalyst 2970 switch. It includes information about VLAN membership modes, VLAN configuration modes, VLAN trunks, and dynamic VLAN assignment from a VLAN Membership Policy Server (VMPS).



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

The chapter includes these sections:

- [Understanding VLANs, page 11-1](#)
- [Configuring Normal-Range VLANs, page 11-4](#)
- [Configuring Extended-Range VLANs, page 11-12](#)
- [Displaying VLANs, page 11-14](#)
- [Configuring VLAN Trunks, page 11-15](#)
- [Configuring VMPS, page 11-26](#)

Understanding VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or a switch supporting fallback bridging, as shown in [Figure 11-1](#). Because a VLAN is considered a separate logical network, it contains its own bridge Management Information Base (MIB) information and can support its own implementation of spanning tree. See [Chapter 14, “Configuring STP.”](#)

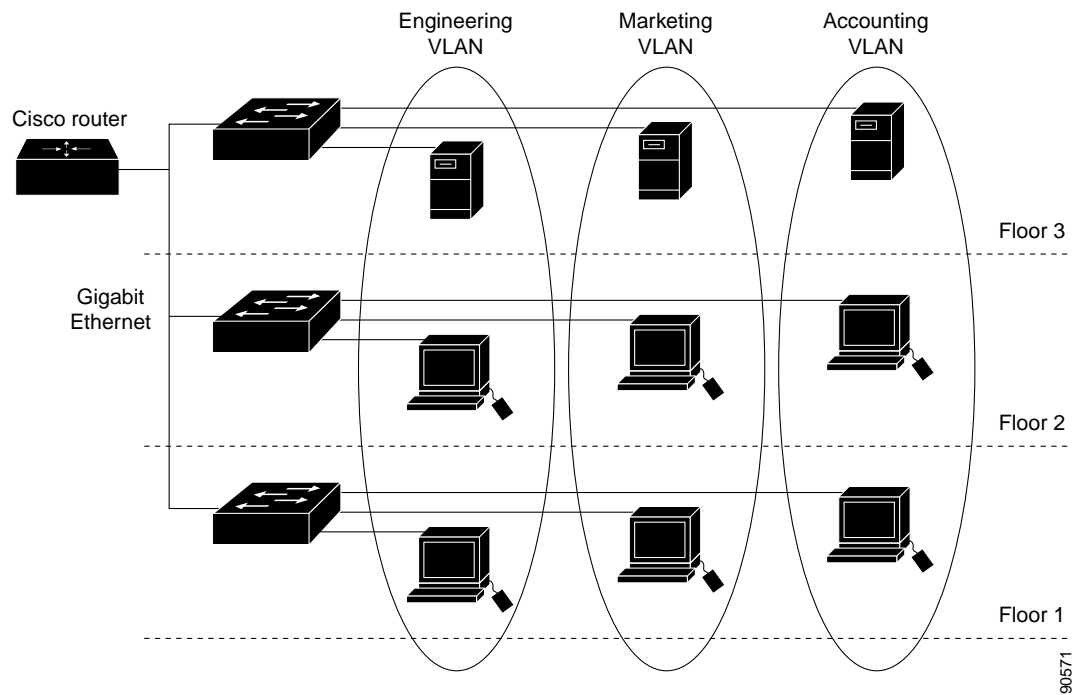


Note

Before you create VLANs, you must decide whether to use VLAN Trunking Protocol (VTP) to maintain global VLAN configuration for your network. For more information on VTP, see [Chapter 12, “Configuring VTP.”](#)

Figure 11-1 shows an example of VLANs segmented into logically defined networks.

Figure 11-1 VLANs as Logically Defined Networks



VLANs are often associated with IP subnetworks. For example, all the end stations in a particular IP subnet belong to the same VLAN. Interface VLAN membership on the switch is assigned manually on an interface-by-interface basis. When you assign switch interfaces to VLANs by using this method, it is known as interface-based, or static, VLAN membership.

Traffic between VLANs must be routed or fallback bridged.

Supported VLANs

The switch supports 1005 VLANs in VTP client, server, and transparent modes. VLANs are identified with a number from 1 to 4094. VLAN IDs 1002 through 1005 are reserved for Token Ring and FDDI VLANs. VTP only learns normal-range VLANs, with VLAN IDs 1 to 1005; VLAN IDs greater than 1005 are extended-range VLANs and are not stored in the VLAN database. The switch must be in VTP transparent mode when you create VLAN IDs from 1006 to 4094.

Although the switch supports a total of 1005 (normal-range and extended-range) VLANs, the number of configured features affects the use of the switch hardware.

The switch supports per-VLAN spanning-tree plus (PVST+) or rapid PVST+ with a maximum of 128 spanning-tree instances. One spanning-tree instance is allowed per VLAN. See the [“Normal-Range VLAN Configuration Guidelines” section on page 11-6](#) for more information about the number of spanning-tree instances and the number of VLANs. The switch supports both Inter-Switch Link (ISL) and IEEE 802.1Q trunking methods for sending VLAN traffic over Ethernet ports.

VLAN Port Membership Modes

You configure a port to belong to a VLAN by assigning a membership mode that determines the kind of traffic the port carries and the number of VLANs to which it can belong. [Table 11-1](#) lists the membership modes and membership and VTP characteristics.

Table 11-1 Port Membership Modes

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Static-access	A static-access port can belong to one VLAN and is manually assigned to that VLAN. For more information, see the “Assigning Static-Access Ports to a VLAN” section on page 11-11 .	VTP is not required. If you do not want VTP to globally propagate information, set the VTP mode to transparent to disable VTP. To participate in VTP, there must be at least one trunk port on the switch connected to a trunk port of a second switch.
Trunk (ISL or IEEE 802.1Q)	A trunk port is a member of all VLANs by default, including extended-range VLANs, but membership can be limited by configuring the allowed-VLAN list. You can also modify the pruning-eligible list to block flooded traffic to VLANs on trunk ports that are included in the list. For information about configuring trunk ports, see the “Configuring an Ethernet Interface as a Trunk Port” section on page 11-18 .	VTP is recommended but not required. VTP maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP exchanges VLAN configuration messages with other switches over trunk links.

Table 11-1 Port Membership Modes (continued)

Membership Mode	VLAN Membership Characteristics	VTP Characteristics
Dynamic access	<p>A dynamic-access port can belong to one VLAN (VLAN ID 1 to 4094) and is dynamically assigned by a VMPS. The VMPS can be a Catalyst 5000 or Catalyst 6000 series switch, for example, but never a Catalyst 2970 switch. The Catalyst 2970 switch is a VMPS client.</p> <p>You can have dynamic-access ports and trunk ports on the same switch, but you must connect the dynamic-access port to an end station or hub and not to another switch.</p> <p>For configuration information, see the “Configuring Dynamic-Access Ports on VMPS Clients” section on page 11-29.</p>	<p>VTP is required.</p> <p>Configure the VMPS and the client with the same VTP domain name.</p> <p>To participate in VTP, there must be at least one trunk port on the switch connected to a trunk port of a second switch.</p>
Voice VLAN	<p>A voice VLAN port is an access port attached to a Cisco IP Phone, configured to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. For more information about voice VLAN ports, see Chapter 13, “Configuring Voice VLAN.”</p>	<p>VTP is not required; it has no affect on voice VLAN.</p>

For more detailed definitions of the modes and their functions, see [Table 11-4 on page 11-16](#).

When a port belongs to a VLAN, the switch learns and manages the addresses associated with the port on a per-VLAN basis. For more information, see the [“Managing the MAC Address Table”](#) section on page 6-22.

Configuring Normal-Range VLANs

Normal-range VLANs are VLANs with VLAN IDs 1 to 1005. If the switch is in VTP server or transparent mode, you can add, modify or remove configurations for VLANs 2 to 1001 in the VLAN database. (VLAN IDs 1 and 1002 to 1005 are automatically created and cannot be removed.)



Note

When the switch is in VTP transparent mode, you can also create extended-range VLANs (VLANs with IDs from 1006 to 4094), but these VLANs are not saved in the VLAN database. See the [“Configuring Extended-Range VLANs”](#) section on page 11-12.

Configurations for VLAN IDs 1 to 1005 are written to the file *vlan.dat* (VLAN database), and you can display them by entering the **show vlan** privileged EXEC command. The *vlan.dat* file is stored in Flash memory.



Caution

You can cause inconsistency in the VLAN database if you attempt to manually delete the *vlan.dat* file. If you want to modify the VLAN configuration, use the commands described in these sections and in the command reference for this release. To change the VTP configuration, see [Chapter 12, “Configuring VTP.”](#)

You use the interface configuration mode to define the port membership mode and to add and remove ports from VLANs. The results of these commands are written to the running-configuration file, and you can display the file by entering the **show running-config** privileged EXEC command.

You can set these parameters when you create a new normal-range VLAN or modify an existing VLAN in the VLAN database:

- VLAN ID
- VLAN name
- VLAN type (Ethernet, Fiber Distributed Data Interface [FDDI], FDDI network entity title [NET], TrBRF, or TrCRF, Token Ring, Token Ring-Net)
- VLAN state (active or suspended)
- Maximum transmission unit (MTU) for the VLAN
- Security Association Identifier (SAID)
- Bridge identification number for TrBRF VLANs
- Ring number for FDDI and TrCRF VLANs
- Parent VLAN number for TrCRF VLANs
- Spanning Tree Protocol (STP) type for TrCRF VLANs
- VLAN number to use when translating from one VLAN type to another

**Note**

This section does not provide configuration details for most of these parameters. For complete information on the commands and parameters that control VLAN configuration, refer to the command reference for this release.

This section includes information about these topics about normal-range VLANs:

- [Token Ring VLANs, page 11-5](#)
- [Normal-Range VLAN Configuration Guidelines, page 11-6](#)
- [VLAN Configuration Mode Options, page 11-6](#)
- [Saving VLAN Configuration, page 11-7](#)
- [Default Ethernet VLAN Configuration, page 11-7](#)
- [Creating or Modifying an Ethernet VLAN, page 11-8](#)
- [Deleting a VLAN, page 11-10](#)
- [Assigning Static-Access Ports to a VLAN, page 11-11](#)

Token Ring VLANs

Although the switch does not support Token Ring connections, a remote device such as a Catalyst 5000 series switch with Token Ring connections could be managed from one of the supported switches. Switches running VTP version 2 advertise information about these Token Ring VLANs:

- Token Ring TrBRF VLANs
- Token Ring TrCRF VLANs

For more information on configuring Token Ring VLANs, refer to the *Catalyst 5000 Series Software Configuration Guide*.

Normal-Range VLAN Configuration Guidelines

Follow these guidelines when creating and modifying normal-range VLANs in your network:

- The switch supports 1005 VLANs in VTP client, server, and transparent modes.
- Normal-range VLANs are identified with a number between 1 and 1001. VLAN numbers 1002 through 1005 are reserved for Token Ring and FDDI VLANs.
- VLAN configuration for VLANs 1 to 1005 are always saved in the VLAN database. If VTP mode is transparent, VTP and VLAN configuration is also saved in the switch running configuration file.
- The switch also supports VLAN IDs 1006 through 4094 in VTP transparent mode (VTP disabled). These are extended-range VLANs and configuration options are limited. Extended-range VLANs are not saved in the VLAN database. See the [“Configuring Extended-Range VLANs” section on page 11-12](#).
- Before you can create a VLAN, the switch must be in VTP server mode or VTP transparent mode. If the switch is a VTP server, you must define a VTP domain or VTP will not function.
- The switch does not support Token Ring or FDDI media. The switch does not forward FDDI, FDDI-Net, TrCRF, or TrBRF traffic, but it does propagate the VLAN configuration through VTP.
- The switch supports 128 spanning-tree instances. If a switch has more active VLANs than supported spanning-tree instances, spanning tree can be enabled on 128 VLANs and is disabled on the remaining VLANs. If you have already used all available spanning-tree instances on a switch, adding another VLAN anywhere in the VTP domain creates a VLAN on that switch that is not running spanning-tree. If you have the default allowed list on the trunk ports of that switch (which is to allow all VLANs), the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that would not be broken, particularly if there are several adjacent switches that all have run out of spanning-tree instances. You can prevent this possibility by setting allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances.

If the number of VLANs on the switch exceeds the number of supported spanning tree instances, we recommend that you configure the IEEE 802.1S Multiple STP (MSTP) on your switch to map multiple VLANs to a single STP instance. For more information about MSTP, see [Chapter 15, “Configuring MSTP.”](#)

VLAN Configuration Mode Options

You can configure normal-range VLANs (with VLAN IDs 1 to 1005) by using these two configuration modes:

- [VLAN Configuration in config-vlan Mode, page 11-6](#)
You access config-vlan mode by entering the **vlan *vlan-id*** global configuration command.
- [VLAN Configuration in VLAN Database Configuration Mode, page 11-7](#)
You access VLAN database configuration mode by entering the **vlan database** privileged EXEC command.

VLAN Configuration in config-vlan Mode

To access config-vlan mode, enter the **vlan** global configuration command with a VLAN ID. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration ([Table 11-2](#)) or enter multiple commands to configure the VLAN. For more

information about commands available in this mode, refer to the **vlan** global configuration command description in the command reference for this release. When you have finished the configuration, you must exit config-vlan mode for the configuration to take effect. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

You must use this config-vlan mode when creating extended-range VLANs (VLAN IDs greater than 1005). See the [“Configuring Extended-Range VLANs” section on page 11-12](#).

VLAN Configuration in VLAN Database Configuration Mode

To access VLAN database configuration mode, enter the **vlan database** privileged EXEC command. Then enter the **vlan** command with a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify the VLAN. You can use the default VLAN configuration ([Table 11-2](#)) or enter multiple commands to configure the VLAN. For more information about keywords available in this mode, refer to the **vlan** VLAN database configuration command description in the command reference for this release. When you have finished the configuration, you must enter **apply** or **exit** for the configuration to take effect. When you enter the **exit** command, it applies all commands and updates the VLAN database. VTP messages are sent to other switches in the VTP domain, and the privileged EXEC mode prompt appears.

Saving VLAN Configuration

The configurations of VLAN IDs 1 to 1005 are always saved in the VLAN database (vlan.dat file). If VTP mode is transparent, they are also saved in the switch running configuration file and you can enter the **copy running-config startup-config** privileged EXEC command to save the configuration in the startup configuration file. To display the VLAN configuration, enter the **show vlan** privileged EXEC command.

When you save VLAN and VTP information (including extended-range VLAN configuration information) in the startup configuration file and reboot the switch, the switch configuration is determined as follows:

- If the VTP mode is transparent in the startup configuration, and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration does not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.
- If VTP mode is server, the domain name and VLAN configuration for the first 1005 VLANs use the VLAN database information



Caution

If the VLAN database configuration is used at startup and the startup configuration file contains extended-range VLAN configuration, this information is lost when the system boots up.

Default Ethernet VLAN Configuration

[Table 11-2](#) shows the default configuration for Ethernet VLANs.

**Note**

The switch supports Ethernet interfaces exclusively. Because FDDI and Token Ring VLANs are not locally supported, you only configure FDDI and Token Ring media-specific characteristics for VTP global advertisements to other switches.

Table 11-2 Ethernet VLAN Defaults and Ranges

Parameter	Default	Range
VLAN ID	1	1 to 4094. Note Extended-range VLANs (VLAN IDs 1006 to 4094) are not saved in the VLAN database.
VLAN name	VLANxxxx, where xxxx represents four numeric digits (including leading zeros) equal to the VLAN ID number	No range
802.10 SAID	100001 (100000 plus the VLAN ID)	1–4294967294
MTU size	1500	1500–18190
Translational bridge 1	0	0–1005
Translational bridge 2	0	0–1005
VLAN state	active	active, suspend
Remote SPAN	disabled	enabled, disabled

Creating or Modifying an Ethernet VLAN

Each Ethernet VLAN in the VLAN database has a unique, 4-digit ID that can be a number from 1 to 1001. VLAN IDs 1002 to 1005 are reserved for Token Ring and FDDI VLANs. To create a normal-range VLAN to be added to the VLAN database, assign a number and name to the VLAN.

**Note**

When the switch is in VTP transparent mode, you can assign VLAN IDs greater than 1006, but they are not added to the VLAN database. See the [“Configuring Extended-Range VLANs” section on page 11-12](#).

For the list of default parameters that are assigned when you add a VLAN, see the [“Configuring Normal-Range VLANs” section on page 11-4](#).

Beginning in privileged EXEC mode, follow these steps to use config-vlan mode to create or modify an Ethernet VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan <i>vlan-id</i>	Enter a VLAN ID, and enter config-vlan mode. Enter a new VLAN ID to create a VLAN, or enter an existing VLAN ID to modify a VLAN. Note The available VLAN ID range for this command is 1 to 4094. For information about adding VLAN IDs greater than 1005 (extended-range VLANs), see the “Configuring Extended-Range VLANs” section on page 11-12.
Step 3	name <i>vlan-name</i>	(Optional) Enter a name for the VLAN. If no name is entered for the VLAN, the default is to append the <i>vlan-id</i> with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
Step 4	mtu <i>mtu-size</i>	(Optional) Change the MTU size (or other VLAN characteristic).
Step 5	remote-span	(Optional) Configure the VLAN as the RSPAN VLAN for a remote SPAN session. For more information on remote SPAN, see Chapter 22, “Configuring SPAN and RSPAN.”
Step 6	end	Return to privileged EXEC mode.
Step 7	show vlan { name <i>vlan-name</i> / id <i>vlan-id</i> }	Verify your entries.
Step 8	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

To return the VLAN name to the default settings, use the **no name**, **no mtu**, or **no remote-span** config-vlan commands.

This example shows how to use config-vlan mode to create Ethernet VLAN 20, name it *test20*, and add it to the VLAN database:

```
Switch# configure terminal
Switch(config)# vlan 20
Switch(config-vlan)# name test20
Switch(config-vlan)# end
```

You can also create or modify Ethernet VLANs by using the VLAN database configuration mode.



Note

VLAN database configuration mode does not support RSPAN VLAN configuration or extended-range VLANs.

Beginning in privileged EXEC mode, follow these steps to use VLAN database configuration mode to create or modify an Ethernet VLAN:

	Command	Purpose
Step 1	vlan database	Enter VLAN database configuration mode.
Step 2	vlan <i>vlan-id</i> name <i>vlan-name</i>	Add an Ethernet VLAN by assigning a number to it. The range is 1 to 1001. You can create or modify a range of consecutive VLANs by entering vlan <i>first-vlan-id</i> end <i>last-vlan-id</i> . Note When entering a VLAN ID in VLAN database configuration mode, do not enter leading zeros. If no name is entered for the VLAN, the default is to append the <i>vlan-id</i> with leading zeros to the word VLAN. For example, VLAN0004 is a default VLAN name for VLAN 4.
Step 3	vlan <i>vlan-id</i> mtu <i>mtu-size</i>	(Optional) To modify a VLAN, identify the VLAN and change a characteristic, such as the MTU size.
Step 4	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 5	show vlan { name <i>vlan-name</i> / id <i>vlan-id</i> }	Verify your entries.
Step 6	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

To return the VLAN name to the default settings, use the **no vlan** *vlan-id* **name** or **no vlan** *vlan-id* **mtu** VLAN database configuration command.

This example shows how to use VLAN configuration mode to create Ethernet VLAN 20, name it *test20*, and add it to the VLAN database:

```
Switch# vlan database
Switch(vlan)# vlan 20 name test20
Switch(vlan)# exit
APPLY completed.
Exiting....
```

Deleting a VLAN

When you delete a VLAN from a switch that is in VTP server mode, the VLAN is removed from the VLAN database for all switches in the VTP domain. When you delete a VLAN from a switch that is in VTP transparent mode, the VLAN is deleted only on that specific switch.

You cannot delete the default VLANs for the different media types: Ethernet VLAN 1 and FDDI or Token Ring VLANs 1002 to 1005.



Caution

When you delete a VLAN, any ports assigned to that VLAN become inactive. They remain associated with the VLAN (and thus inactive) until you assign them to a new VLAN.

Beginning in privileged EXEC mode, follow these steps to delete a VLAN on the switch by using global configuration mode:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no vlan <i>vlan-id</i>	Remove the VLAN by entering the VLAN ID.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vlan brief	Verify the VLAN removal.
Step 5	copy running-config startup config	(Optional) If the switch is in VTP transparent mode, the VLAN configuration is saved in the running configuration file as well as in the VLAN database. This saves the configuration in the switch startup configuration file.

To delete a VLAN by using VLAN database configuration mode, use the **vlan database** privileged EXEC command to enter VLAN database configuration mode and the **no vlan** *vlan-id* VLAN database configuration command.

Assigning Static-Access Ports to a VLAN

You can assign a static-access port to a VLAN without having VTP globally propagate VLAN configuration information by disabling VTP (VTP transparent mode).

If you are assigning a port on a cluster member switch to a VLAN, first use the **rcommand** privileged EXEC command to log in to the cluster member switch.



Note

If you assign an interface to a VLAN that does not exist, the new VLAN is created. (See the [“Creating or Modifying an Ethernet VLAN”](#) section on page 11-8.)

Beginning in privileged EXEC mode, follow these steps to assign a port to a VLAN in the VLAN database:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	interface <i>interface-id</i>	Enter the interface to be added to the VLAN.
Step 3	switchport mode access	Define the VLAN membership mode for the port (Layer 2 access port).
Step 4	switchport access vlan <i>vlan-id</i>	Assign the port to a VLAN. Valid VLAN IDs are 1 to 4094.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config interface <i>interface-id</i>	Verify the VLAN membership mode of the interface.
Step 7	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Administrative Mode</i> and the <i>Access Mode VLAN</i> fields of the display.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/1 as an access port in VLAN 2:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport access vlan 2
Switch(config-if)# end
```

Configuring Extended-Range VLANs

When the switch is in VTP transparent mode (VTP disabled), you can create extended-range VLANs (in the range 1006 to 4094). Extended-range VLANs enable service providers to extend their infrastructure to a greater number of customers. The extended-range VLAN IDs are allowed for any switchport commands that allow VLAN IDs. You always use config-vlan mode (accessed by entering the **vlan** *vlan-id* global configuration command) to configure extended-range VLANs. The extended range is not supported in VLAN database configuration mode (accessed by entering the **vlan database** privileged EXEC command).

Extended-range VLAN configurations are not stored in the VLAN database, but because VTP mode is transparent, they are stored in the switch running configuration file, and you can save the configuration in the startup configuration file by using the **copy running-config startup-config** privileged EXEC command.



Note

Although the switch supports 4094 VLAN IDs, see the [“Supported VLANs” section on page 11-3](#) for the actual number of VLANs supported.

This section includes this information about extended-range VLANs:

- [Default VLAN Configuration, page 11-12](#)
- [Extended-Range VLAN Configuration Guidelines, page 11-12](#)
- [Creating an Extended-Range VLAN, page 11-13](#)

Default VLAN Configuration

See [Table 11-2 on page 11-8](#) for the default configuration for Ethernet VLANs. You can change only the MTU size and remote SPAN configuration state on extended-range VLANs; all other characteristics must remain at the default state.

Extended-Range VLAN Configuration Guidelines

Follow these guidelines when creating extended-range VLANs:

- To add an extended-range VLAN, you must use the **vlan** *vlan-id* global configuration command and access config-vlan mode. You cannot add extended-range VLANs in VLAN database configuration mode (accessed by entering the **vlan database** privileged EXEC command).

- VLAN IDs in the extended range are not saved in the VLAN database and are not recognized by VTP.
- You cannot include extended-range VLANs in the pruning eligible range.
- The switch must be in VTP transparent mode when you create extended-range VLANs. If VTP mode is server or client, an error message is generated, and the extended-range VLAN is rejected.
- You can set the VTP mode to transparent in global configuration mode or in VLAN database configuration mode. See the [“Disabling VTP \(VTP Transparent Mode\)” section on page 12-11](#). You should save this configuration to the startup configuration so that the switch boots up in VTP transparent mode. Otherwise, you lose the extended-range VLAN configuration if the switch resets.
- STP is enabled by default on extended-range VLANs, but you can disable it by using the **no spanning-tree vlan *vlan-id*** global configuration command. When the maximum number of spanning-tree instances (128) are on the switch, spanning tree is disabled on any newly created VLANs. If the number of VLANs on the switch exceeds the maximum number of spanning tree instances, we recommend that you configure the IEEE 802.1S Multiple STP (MSTP) on your switch to map multiple VLANs to a single STP instance. For more information about MSTP, see [Chapter 15, “Configuring MSTP.”](#)
- Although the switch supports a total of 1005 (normal-range and extended-range) VLANs, the number of configured features affects the use of the switch hardware. If you try to create an extended-range VLAN and there are not enough hardware resources available, an error message is generated, and the extended-range VLAN is rejected.

Creating an Extended-Range VLAN

You create an extended-range VLAN in global configuration mode by entering the **vlan** global configuration command with a VLAN ID from 1006 to 4094. This command accesses the config-vlan mode. The extended-range VLAN has the default Ethernet VLAN characteristics (see [Table 11-2](#)) and the MTU size and RSPAN configuration are the only parameters you can change. Refer to the description of the **vlan** global configuration command in the command reference for defaults of all parameters. If you enter an extended-range VLAN ID when the switch is not in VTP transparent mode, an error message is generated when you exit from config-vlan mode, and the extended-range VLAN is not created.

Extended-range VLANs are not saved in the VLAN database; they are saved in the switch running configuration file. You can save the extended-range VLAN configuration in the switch startup configuration file by using the **copy running-config startup-config** privileged EXEC command.

Beginning in privileged EXEC mode, follow these steps to create an extended-range VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode transparent	Configure the switch for VTP transparent mode, disabling VTP.
Step 3	vlan <i>vlan-id</i>	Enter an extended-range VLAN ID and enter config-vlan mode. The range is 1006 to 4094.
Step 4	mtu <i>mtu-size</i>	(Optional) Modify the VLAN by changing the MTU size. Note Although all VLAN commands appear in the CLI help in config-vlan mode, only the mtu <i>mtu-size</i> and remote-span commands are supported for extended-range VLANs.

	Command	Purpose
Step 5	remote-span	(Optional) Configure the VLAN as the RSPAN VLAN. See the “Configuring a VLAN as an RSPAN VLAN” section on page 22-17.
Step 6	end	Return to privileged EXEC mode.
Step 7	show vlan id <i>vlan-id</i>	Verify that the VLAN has been created.
Step 8	copy running-config startup config	Save your entries in the switch startup configuration file. To save extended-range VLAN configurations, you need to save the VTP transparent mode configuration and the extended-range VLAN configuration in the switch startup configuration file. Otherwise, if the switch resets, it will default to VTP server mode, and the extended-range VLAN IDs will not be saved.

To delete an extended-range VLAN, use the **no vlan** *vlan-id* global configuration command.

The procedure for assigning static-access ports to an extended-range VLAN is the same as for normal-range VLANs. See the [“Assigning Static-Access Ports to a VLAN”](#) section on page 11-11.

This example shows how to create a new extended-range VLAN with all default characteristics, enter config-vlan mode, and save the new VLAN in the switch startup configuration file:

```
Switch(config)# vtp mode transparent
Switch(config)# vlan 2000
Switch(config-vlan)# end
Switch# copy running-config startup config
```

Displaying VLANs

Use the **show vlan** privileged EXEC command to display a list of all VLANs on the switch, including extended-range VLANs. The display includes VLAN status, ports, and configuration information. To view normal-range VLANs in the VLAN database (1 to 1005), use the **show** VLAN database configuration command (accessed by entering the **vlan database** privileged EXEC command).

[Table 11-3](#) lists the commands for monitoring VLANs.

Table 11-3 VLAN Monitoring Commands

Command	Command Mode	Purpose
show	VLAN database configuration	Display status of VLANs in the VLAN database.
show current [<i>vlan-id</i>]	VLAN database configuration	Display status of all or the specified VLAN in the VLAN database.
show interfaces [vlan <i>vlan-id</i>]	Privileged EXEC	Display characteristics for all interfaces or for the specified VLAN configured on the switch.
show vlan [id <i>vlan-id</i>]	Privileged EXEC	Display parameters for all VLANs or the specified VLAN on the switch.

For more details about the show command options and explanations of output fields, refer to the command reference for this release.

Configuring VLAN Trunks

These sections describe how VLAN trunks function on the switch:

- [Trunking Overview, page 11-15](#)
- [Encapsulation Types, page 11-16](#)
- [Default Layer 2 Ethernet Interface VLAN Configuration, page 11-17](#)
- [Configuring an Ethernet Interface as a Trunk Port, page 11-18](#)
- [Configuring Trunk Ports for Load Sharing, page 11-22](#)

Trunking Overview

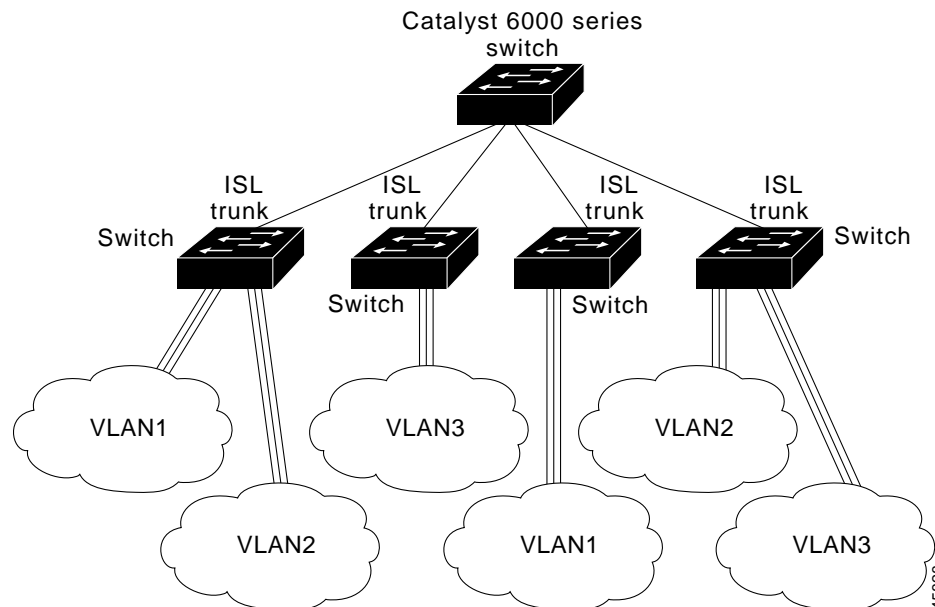
A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network.

Two trunking encapsulations are available on all Ethernet interfaces:

- Inter-Switch Link (ISL)—ISL is Cisco-proprietary trunking encapsulation.
- 802.1Q—802.1Q is industry-standard trunking encapsulation.

[Figure 11-2](#) shows a network of switches that are connected by ISL trunks.

Figure 11-2 *Switches in an ISL Trunking Environment*



You can configure a trunk on a single Ethernet interface or on an EtherChannel bundle. For more information about EtherChannel, see [Chapter 28, “Configuring EtherChannels.”](#)

Ethernet trunk interfaces support different trunking modes (see [Table 11-4](#)). You can set an interface as trunking or nontrunking or to negotiate trunking with the neighboring interface. To autonegotiate trunking, the interfaces must be in the same VTP domain.

Trunk negotiation is managed by the Dynamic Trunking Protocol (DTP), which is a Point-to-Point Protocol. However, some internetworking devices might forward DTP frames improperly, which could cause misconfigurations.

To avoid this, you should configure interfaces connected to devices that do not support DTP to not forward DTP frames, that is, to turn off DTP.

- If you do not intend to trunk across those links, use the **switchport mode access** interface configuration command to disable trunking.
- To enable trunking to a device that does not support DTP, use the **switchport mode trunk** and **switchport nonegotiate** interface configuration commands to cause the interface to become a trunk but to not generate DTP frames. Use the **switchport trunk encapsulation isl** or **switchport trunk encapsulation dot1q** interface to select the encapsulation type on the trunk port.

You can also specify on DTP interfaces whether the trunk uses ISL or 802.1Q encapsulation or if the encapsulation type is autonegotiated. The DTP supports autonegotiation of both ISL and 802.1Q trunks.

Table 11-4 Layer 2 Interface Modes

Mode	Function
switchport mode access	Puts the interface (access port) into permanent nontrunking mode and negotiates to convert the link into a nontrunk link. The interface becomes a nontrunk interface regardless of whether or not the neighboring interface is a trunk interface.
switchport mode dynamic auto	Makes the interface able to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to <i>trunk</i> or <i>desirable</i> mode. The default switchport mode for all Ethernet interfaces is dynamic auto .
switchport mode dynamic desirable	Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to <i>trunk</i> , <i>desirable</i> , or <i>auto</i> mode.
switchport mode trunk	Puts the interface into permanent trunking mode and negotiates to convert the neighboring link into a trunk link. The interface becomes a trunk interface even if the neighboring interface is not a trunk interface.
switchport nonegotiate	Prevents the interface from generating DTP frames. You can use this command only when the interface switchport mode is access or trunk . You must manually configure the neighboring interface as a trunk interface to establish a trunk link.

Encapsulation Types

Table 11-5 lists the Ethernet trunk encapsulation types and keywords.

Table 11-5 Ethernet Trunk Encapsulation Types

Encapsulation	Function
switchport trunk encapsulation isl	Specifies ISL encapsulation on the trunk link.
switchport trunk encapsulation dot1q	Specifies 802.1Q encapsulation on the trunk link.
switchport trunk encapsulation negotiate	Specifies that the interface negotiate with the neighboring interface to become an ISL (preferred) or 802.1Q trunk, depending on the configuration and capabilities of the neighboring interface. This is the default for the switch.

The trunking mode, the trunk encapsulation type, and the hardware capabilities of the two connected interfaces determine whether a link becomes an ISL or 802.1Q trunk.

802.1Q Configuration Considerations

802.1Q trunks impose these limitations on the trunking strategy for a network:

- In a network of Cisco switches connected through 802.1Q trunks, the switches maintain one instance of spanning tree for each VLAN allowed on the trunks. Non-Cisco devices might support one spanning-tree instance for all VLANs.

When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch combines the spanning-tree instance of the VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch. However, spanning-tree information for each VLAN is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

- Make sure the native VLAN for an 802.1Q trunk is the same on both ends of the trunk link. If the native VLAN on one end of the trunk is different from the native VLAN on the other end, spanning-tree loops might result.
- Disabling spanning tree on the native VLAN of an 802.1Q trunk without disabling spanning tree on every VLAN in the network can potentially cause spanning-tree loops. We recommend that you leave spanning tree enabled on the native VLAN of an 802.1Q trunk or disable spanning tree on every VLAN in the network. Make sure your network is loop-free before disabling spanning tree.

Default Layer 2 Ethernet Interface VLAN Configuration

Table 11-6 shows the default Layer 2 Ethernet interface VLAN configuration.

Table 11-6 Default Layer 2 Ethernet Interface VLAN Configuration

Feature	Default Setting
Interface mode	switchport mode dynamic auto
Trunk encapsulation	switchport trunk encapsulation negotiate
Allowed VLAN range	VLANs 1 to 4094
VLAN range eligible for pruning	VLANs 2 to 1001
Default VLAN (for access ports)	VLAN 1
Native VLAN (for 802.1Q trunks)	VLAN 1

Configuring an Ethernet Interface as a Trunk Port

Because trunk ports send and receive VTP advertisements, to use VTP you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of a second switch. Otherwise, the switch cannot receive any VTP advertisements.

This section includes these procedures for configuring an Ethernet interface as a trunk port on the switch:

- [Interaction with Other Features, page 11-18](#)
- [Defining the Allowed VLANs on a Trunk, page 11-20](#)
- [Changing the Pruning-Eligible List, page 11-21](#)
- [Configuring the Native VLAN for Untagged Traffic, page 11-21](#)

**Note**

By default, trunks negotiate encapsulation. If the neighboring interface supports ISL and 802.1Q encapsulation and both interfaces are set to negotiate the encapsulation type, the trunk uses ISL encapsulation.

Interaction with Other Features

Trunking interacts with other features in these ways:

- A trunk port cannot be a secure port.
- Trunk ports can be grouped into EtherChannel port groups, but all trunks in the group must have the same configuration. When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, the switch propagates the setting you entered to all ports in the group:
 - allowed-VLAN list
 - STP port priority for each VLAN
 - STP Port Fast setting
 - trunk status: if one port in a port group ceases to be a trunk, all ports cease to be trunks.
- We recommend that you configure no more than 24 trunk ports in PVST mode and no more than 40 trunk ports in MST mode.
- If you try to enable 802.1X on a trunk port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to trunk, the port mode is not changed.
- A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1X on a dynamic port, an error message appears, and 802.1X is not enabled. If you try to change the mode of an 802.1X-enabled port to dynamic, the port mode is not changed.

Configuring a Trunk Port

Beginning in privileged EXEC mode, follow these steps to configure a port as an ISL or 802.1Q trunk port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter the interface configuration mode and the port to be configured for trunking.
Step 3	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate (the default) with the neighboring interface for encapsulation type. You must configure each end of the link with the same encapsulation type.
Step 4	switchport mode {dynamic {auto desirable} trunk}	Configure the interface as a Layer 2 trunk (required only if the interface is a Layer 2 access port or to specify the trunking mode). <ul style="list-style-type: none"> dynamic auto—Set the interface to a trunk link if the neighboring interface is set to trunk or desirable mode. This is the default. dynamic desirable—Set the interface to a trunk link if the neighboring interface is set to trunk, desirable, or auto mode. trunk—Set the interface in permanent trunking mode and negotiate to convert the link to a trunk link even if the neighboring interface is not a trunk interface.
Step 5	switchport access vlan <i>vlan-id</i>	(Optional) Specify the default VLAN, which is used if the interface stops trunking.
Step 6	switchport trunk native vlan <i>vlan-id</i>	Specify the native VLAN for 802.1Q trunks.
Step 7	end	Return to privileged EXEC mode.
Step 8	show interfaces <i>interface-id</i> switchport	Display the switchport configuration of the interface in the <i>Administrative Mode</i> and the <i>Administrative Trunking Encapsulation</i> fields of the display.
Step 9	show interfaces <i>interface-id</i> trunk	Display the trunk configuration of the interface.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To reset all trunking characteristics of a trunking interface to the defaults, use the **no switchport trunk** interface configuration command. To disable trunking, use the **switchport mode access** interface configuration command to configure the port as a static-access port.

This example shows how to configure the Gigabit Ethernet interface 0/4 as an 802.1Q trunk. The example assumes that the neighbor interface is configured to support 802.1Q trunking.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# switchport mode dynamic desirable
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# end
```

Defining the Allowed VLANs on a Trunk

By default, a trunk port sends traffic to and receives traffic from all VLANs. All VLAN IDs, 1 to 4094, are allowed on each trunk. However, you can remove VLANs from the allowed list, preventing traffic from those VLANs from passing over the trunk. To restrict the traffic a trunk carries, use the **switchport trunk allowed vlan remove *vlan-list*** interface configuration command to remove specific VLANs from the allowed list.



Note

VLAN 1 is the default VLAN on all trunk ports in all Cisco switches, and it has previously been a requirement that VLAN 1 always be enabled on every trunk link. You can use the VLAN 1 minimization feature to disable VLAN 1 on any individual VLAN trunk link so that no user traffic (including spanning tree advertisements) is sent or received on VLAN 1.

To reduce the risk of spanning-tree loops or storms, you can disable VLAN 1 on any individual VLAN trunk port by removing VLAN 1 from the allowed list. When you remove VLAN 1 from a trunk port, the interface continues to send and receive management traffic, for example, Cisco Discovery Protocol (CDP), Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP), Dynamic Trunking Protocol (DTP), and VLAN Trunking Protocol (VTP) in VLAN 1.

If a trunk port with VLAN 1 disabled is converted to a nontrunk port, it is added to the access VLAN. If the access VLAN is set to 1, the port will be added to VLAN 1, regardless of the **switchport trunk allowed** setting. The same is true for any VLAN that has been disabled on the port.

A trunk port can become a member of a VLAN if the VLAN is enabled, if VTP knows of the VLAN, and if the VLAN is in the allowed list for the port. When VTP detects a newly enabled VLAN and the VLAN is in the allowed list for a trunk port, the trunk port automatically becomes a member of the enabled VLAN. When VTP detects a new VLAN and the VLAN is not in the allowed list for a trunk port, the trunk port does not become a member of the new VLAN.

Beginning in privileged EXEC mode, follow these steps to modify the allowed list of an ISL or 802.1Q trunk:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode and the port to be configured.
Step 3	switchport mode trunk	Configure the interface as a VLAN trunk port.
Step 4	switchport trunk allowed vlan {add all except remove} <i>vlan-list</i>	<p>(Optional) Configure the list of VLANs allowed on the trunk.</p> <p>For explanations about using the add, all, except, and remove keywords, refer to the command reference for this release.</p> <p>The <i>vlan-list</i> parameter is either a single VLAN number from 1 to 4094 or a range of VLANs described by two VLAN numbers, the lower one first, separated by a hyphen. Do not enter any spaces between comma-separated VLAN parameters or in hyphen-specified ranges.</p> <p>All VLANs are allowed by default.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking VLANs Enabled</i> field of the display.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default allowed VLAN list of all VLANs, use the **no switchport trunk allowed vlan** interface configuration command.

This example shows how to remove VLAN 2 from the allowed VLAN list on an interface:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport trunk allowed vlan remove 2
Switch(config-if)# end
```

Changing the Pruning-Eligible List

The pruning-eligible list applies only to trunk ports. Each trunk port has its own eligibility list. VTP pruning must be enabled for this procedure to take effect. The [“Enabling VTP Pruning” section on page 12-13](#) describes how to enable VTP pruning.

Beginning in privileged EXEC mode, follow these steps to remove VLANs from the pruning-eligible list on a trunk port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and select the trunk port for which VLANs should be pruned.
Step 3	switchport trunk pruning vlan { add except none remove } <i>vlan-list</i> [, <i>vlan</i> [, <i>vlan</i> [,...]]]	Configure the list of VLANs allowed to be pruned from the trunk. (See the “VTP Pruning” section on page 12-4). For explanations about using the add , except , none , and remove keywords, refer to the command reference for this release. Separate nonconsecutive VLAN IDs with a comma and no spaces; use a hyphen to designate a range of IDs. Valid IDs are from 2 to 1001. Extended-range VLANs (VLAN IDs 1006 to 4094) cannot be pruned. VLANs that are pruning-ineligible receive flooded traffic. The default list of VLANs allowed to be pruned contains VLANs 2 to 1001.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Pruning VLANs Enabled</i> field of the display.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default pruning-eligible list of all VLANs, use the **no switchport trunk pruning vlan** interface configuration command.

Configuring the Native VLAN for Untagged Traffic

A trunk port configured with 802.1Q tagging can receive both tagged and untagged traffic. By default, the switch forwards untagged traffic in the native VLAN configured for the port. The native VLAN is VLAN 1 by default.



Note

The native VLAN can be assigned any VLAN ID.

For information about 802.1Q configuration issues, see the [“802.1Q Configuration Considerations” section on page 11-17](#).

Beginning in privileged EXEC mode, follow these steps to configure the native VLAN on an 802.1Q trunk:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and define the interface that is configured as the 802.1Q trunk.
Step 3	switchport trunk native vlan <i>vlan-id</i>	Configure the VLAN that is sending and receiving untagged traffic on the trunk port. For <i>vlan-id</i> , the range is 1 to 4094.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Trunking Native Mode VLAN</i> field.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default native VLAN, VLAN 1, use the **no switchport trunk native vlan** interface configuration command.

If a packet has a VLAN ID that is the same as the outgoing port native VLAN ID, the packet is sent untagged; otherwise, the switch sends the packet with a tag.

Configuring Trunk Ports for Load Sharing

Load sharing divides the bandwidth supplied by parallel trunks connecting switches. To avoid loops, STP normally blocks all but one parallel link between switches. Using load sharing, you divide the traffic between the links according to which VLAN the traffic belongs.

You configure load sharing on trunk ports by using STP port priorities or STP path costs. For load sharing using STP port priorities, both load-sharing links must be connected to the same switch. For load sharing using STP path costs, each load-sharing link can be connected to the same switch or to two different switches. For more information about STP, see [Chapter 14, “Configuring STP.”](#)

Load Sharing Using STP Port Priorities

When two ports on the same switch form a loop, the STP port priority setting determines which port is enabled and which port is in a blocking state. You can set the priorities on a parallel trunk port so that the port carries all the traffic for a given VLAN. The trunk port with the higher priority (lower values) for a VLAN is forwarding traffic for that VLAN. The trunk port with the lower priority (higher values) for the same VLAN remains in a blocking state for that VLAN. One trunk port sends or receives all traffic for the VLAN.

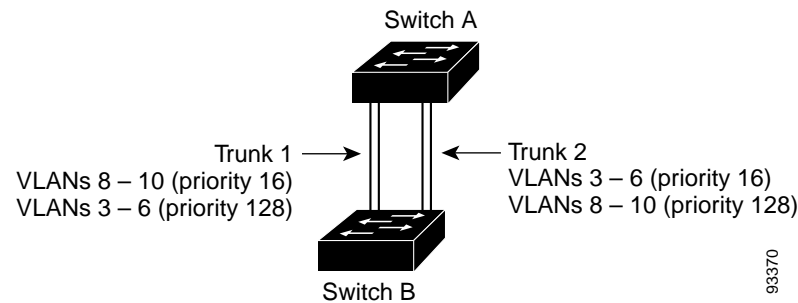
[Figure 11-3](#) shows two trunks connecting supported switches. In this example, the switches are configured as follows:

- VLANs 8 through 10 are assigned a port priority of 16 on Trunk 1.
- VLANs 3 through 6 retain the default port priority of 128 on Trunk 1.
- VLANs 3 through 6 are assigned a port priority of 16 on Trunk 2.

- VLANs 8 through 10 retain the default port priority of 128 on Trunk 2.

In this way, Trunk 1 carries traffic for VLANs 8 through 10, and Trunk 2 carries traffic for VLANs 3 through 6. If the active trunk fails, the trunk with the lower priority takes over and carries the traffic for all of the VLANs. No duplication of traffic occurs over any trunk port.

Figure 11-3 Load Sharing by Using STP Port Priorities



Beginning in privileged EXEC mode, follow these steps to configure the network shown in [Figure 11-3](#).

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode on Switch A.
Step 2	vtp domain <i>domain-name</i>	Configure a VTP administrative domain. The domain name can be from 1 to 32 characters.
Step 3	vtp mode server	Configure Switch A as the VTP server.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vtp status	Verify the VTP configuration on both Switch A and Switch B. In the display, check the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields.
Step 6	show vlan	Verify that the VLANs exist in the database on Switch A.
Step 7	configure terminal	Enter global configuration mode.
Step 8	interface gigabitethernet 0/1	Enter interface configuration mode, and define Gigabit Ethernet port 0/1 as the interface to be configured as a trunk.
Step 9	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation or to negotiate with the neighboring interface. You must configure each end of the link with the same encapsulation type.
Step 10	switchport mode trunk	Configure the port as a trunk port.
Step 11	end	Return to privileged EXEC mode.
Step 12	show interfaces gigabitethernet 0/1 switchport	Verify the VLAN configuration.
Step 13		Repeat Steps 7 through 11 on Switch A for a second interface in the switch.
Step 14		Repeat Steps 7 through 11 on Switch B to configure the trunk ports that connect to the trunk ports configured on Switch A.

	Command	Purpose
Step 15	show vlan	When the trunk links come up, VTP passes the VTP and VLAN information to Switch B. Verify that Switch B has learned the VLAN configuration.
Step 16	configure terminal	Enter global configuration mode on Switch A.
Step 17	interface gigabitethernet 0/1	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 18	spanning-tree vlan 8-10 port-priority 16	Assign the port priority of 10 for VLANs 8 through 10.
Step 19	exit	Return to global configuration mode.
Step 20	interface gigabitethernet0/2	Enter interface configuration mode, and define the interface to set the STP port priority.
Step 21	spanning-tree vlan 3-6 port-priority 16	Assign the port priority of 10 for VLANs 3 through 6.
Step 22	end	Return to privileged EXEC mode.
Step 23	show running-config	Verify your entries.
Step 24	copy running-config startup-config	(Optional) Save your entries in the configuration file.

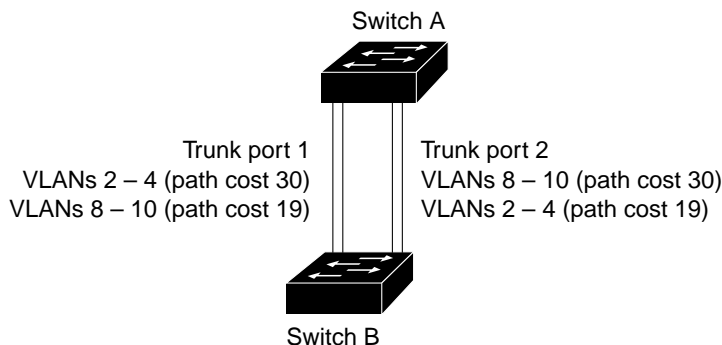
Load Sharing Using STP Path Cost

You can configure parallel trunks to share VLAN traffic by setting different path costs on a trunk and associating the path costs with different sets of VLANs, blocking different ports for different VLANs. The VLANs keep the traffic separate and maintain redundancy in the event of a lost link.

In [Figure 11-4](#), Trunk ports 1 and 2 are configured as 100BASE-T ports. These VLAN path costs are assigned:

- VLANs 2 through 4 are assigned a path cost of 30 on Trunk port 1.
- VLANs 8 through 10 retain the default 100BASE-T path cost on Trunk port 1 of 19.
- VLANs 8 through 10 are assigned a path cost of 30 on Trunk port 2.
- VLANs 2 through 4 retain the default 100BASE-T path cost on Trunk port 2 of 19.

Figure 11-4 Load-Sharing Trunks with Traffic Distributed by Path Cost



90573

Beginning in privileged EXEC mode, follow these steps to configure the network shown in [Figure 11-4](#):

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode on Switch A.
Step 2	interface gigabitethernet0/1	Enter interface configuration mode, and define Gigabit Ethernet port 0/1 as the interface to be configured as a trunk.
Step 3	switchport trunk encapsulation {isl dot1q negotiate}	Configure the port to support ISL or 802.1Q encapsulation. You must configure each end of the link with the same encapsulation type.
Step 4	switchport mode trunk	Configure the port as a trunk port. The trunk defaults to ISL trunking.
Step 5	exit	Return to global configuration mode.
Step 6		Repeat Steps 2 through 4 on a second interface in Switch A.
Step 7	end	Return to privileged EXEC mode.
Step 8	show running-config	Verify your entries. In the display, make sure that the interfaces configured in Steps 2 and 6 are configured as trunk ports.
Step 9	show vlan	When the trunk links come up, Switch A receives the VTP information from the other switches. Verify that Switch A has learned the VLAN configuration.
Step 10	configure terminal	Enter global configuration mode.
Step 11	interface gigabitethernet0/1	Enter interface configuration mode, and define Gigabit Ethernet port 0/1 as the interface on which to set the STP cost.
Step 12	spanning-tree vlan 2-4 cost 30	Set the spanning-tree path cost to 30 for VLANs 2 through 4.
Step 13	end	Return to global configuration mode.
Step 14		Repeat Steps 9 through 11 on the other configured trunk interface on Switch A, and set the spanning-tree path cost to 30 for VLANs 8, 9, and 10.
Step 15	exit	Return to privileged EXEC mode.
Step 16	show running-config	Verify your entries. In the display, verify that the path costs are set correctly for both trunk interfaces.
Step 17	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring VMPS

The VLAN Query Protocol (VQP) is used to support dynamic-access ports, which are not permanently assigned to a VLAN, but given VLAN assignments based on the MAC source addresses seen on the port. Each time an unknown MAC address is seen, the switch sends a VQP query to a remote VMPS; the query includes the newly seen MAC address and the port on which it was seen. The VMPS responds with a VLAN assignment for the port. The switch cannot be a VMPS server but can act as a client to the VMPS and communicate with it through VQP.

This section includes this information about configuring VMPS:

- [“Understanding VMPS” section on page 11-26](#)
- [“Default VMPS Client Configuration” section on page 11-27](#)
- [“VMPS Configuration Guidelines” section on page 11-27](#)
- [“Configuring the VMPS Client” section on page 11-28](#)
- [“Monitoring the VMPS” section on page 11-30](#)
- [“Troubleshooting Dynamic-Access Port VLAN Membership” section on page 11-31](#)
- [“VMPS Configuration Example” section on page 11-31](#)

Understanding VMPS

Each time the client switch receives the MAC address of a new host, it sends a VQP query to the VMPS. When the VMPS receives this query, it searches its database for a MAC-address-to-VLAN mapping. The server response is based on this mapping and whether or not the server is in open or secure mode. In secure mode, the server shuts down the port when an illegal host is detected. In open mode, the server simply denies the host access to the port.

If the port is currently *unassigned* (that is, it does not yet have a VLAN assignment), the VMPS provides one of these responses:

- If the host is allowed on the port, the VMPS sends the client a *vlan-assignment* response containing the assigned VLAN name and allowing access to the host.
- If the host is not allowed on the port and the VMPS is in open mode, the VMPS sends an *access-denied* response.
- If the VLAN is not allowed on the port and the VMPS is in secure mode, the VMPS sends a *port-shutdown* response.

If the port already has a VLAN assignment, the VMPS provides one of these responses:

- If the VLAN in the database matches the current VLAN on the port, the VMPS sends an *success* response, allowing access to the host.
- If the VLAN in the database does not match the current VLAN on the port and active hosts exist on the port, the VMPS sends an *access-denied* or a *port-shutdown* response, depending on the secure mode of the VMPS.

If the switch receives an *access-denied* response from the VMPS, it continues to block traffic to and from the host MAC address. The switch continues to monitor the packets directed to the port and sends a query to the VMPS when it identifies a new host address. If the switch receives a *port-shutdown* response from the VMPS, it disables the port. The port must be manually re-enabled by using the CLI, CMS, or SNMP.

Dynamic-Access Port VLAN Membership

A dynamic-access port can belong to only one VLAN with an ID from 1 to 4094. When the link comes up, the switch does not forward traffic to or from this port until the VMPS provides the VLAN assignment. The VMPS receives the source MAC address from the first packet of a new host connected to the dynamic-access port and attempts to match the MAC address to a VLAN in the VMPS database.

If there is a match, the VMPS sends the VLAN number for that port. If the client switch was not previously configured, it uses the domain name from the first VTP packet it receives on its trunk port from the VMPS. If the client switch was previously configured, it includes its domain name in the query packet to the VMPS to obtain its VLAN number. The VMPS verifies that the domain name in the packet matches its own domain name before accepting the request and responds to the client with the assigned VLAN number for the client. If there is no match, the VMPS either denies the request or shuts down the port (depending on the VMPS secure mode setting).

Multiple hosts (MAC addresses) can be active on a dynamic-access port if they are all in the same VLAN; however, the VMPS shuts down a dynamic-access port if more than 20 hosts are active on the port.

If the link goes down on a dynamic-access port, the port returns to an isolated state and does not belong to a VLAN. Any hosts that come online through the port are checked again through the VQP with the VMPS before the port is assigned to a VLAN.

Dynamic-access ports can be used for direct host connections, or they can connect to a network. A maximum of 20 MAC addresses are allowed per port on the switch. A dynamic-access port can belong to only one VLAN at a time, but the VLAN can change over time, depending on the MAC addresses seen.

Default VMPS Client Configuration

Table 11-7 shows the default VMPS and dynamic-access port configuration on client switches.

Table 11-7 Default VMPS Client and Dynamic-Access Port Configuration

Feature	Default Setting
VMPS domain server	None
VMPS reconfirm interval	60 minutes
VMPS server retry count	3
Dynamic-access ports	None configured

VMPS Configuration Guidelines

These guidelines and restrictions apply to dynamic-access port VLAN membership:

- You should configure the VMPS before you configure ports as dynamic-access ports.
- When you configure a port as a dynamic-access port, the spanning-tree Port Fast feature is automatically enabled for that port. The Port Fast mode accelerates the process of bringing the port into the forwarding state.

- 802.1X ports cannot be configured as dynamic-access ports. If you try to enable 802.1X on a dynamic-access (VQP) port, an error message appears, and 802.1X is not enabled. If you try to change an 802.1X-enabled port to dynamic VLAN assignment, an error message appears, and the VLAN configuration is not changed.
- Trunk ports cannot be dynamic-access ports, but you can enter the **switchport access vlan dynamic** interface configuration command for a trunk port. In this case, the switch retains the setting and applies it if the port is later configured as an access port.
You must turn off trunking on the port before the dynamic-access setting takes effect.
- Dynamic-access ports cannot be monitor ports.
- Secure ports cannot be dynamic-access ports. You must disable port security on a port before it becomes dynamic.
- Dynamic-access ports cannot be members of an EtherChannel group.
- Port channels cannot be configured as dynamic-access ports.
- The VTP management domain of the VMPS client and the VMPS server must be the same.
- The VLAN configured on the VMPS server should not be a voice VLAN.

Configuring the VMPS Client

You configure dynamic VLANs by using the VMPS (server). The switch can be a VMPS client; it cannot be a VMPS server.

Entering the IP Address of the VMPS

You must first enter the IP address of the server to configure the switch as a client.



Note

If the VMPS is being defined for a cluster of switches, enter the address on the command switch.

Beginning in privileged EXEC mode, follow these steps to enter the IP address of the VMPS:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps server <i>ipaddress</i> primary	Enter the IP address of the switch acting as the primary VMPS server.
Step 3	vmps server <i>ipaddress</i>	(Optional) Enter the IP address of the switch acting as a secondary VMPS server. You can enter up to three secondary server addresses.
Step 4	end	Return to privileged EXEC mode.
Step 5	show vmps	Verify your entries in the <i>VMPS Domain Server</i> field of the display.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

You must have IP connectivity to the VMPS for dynamic-access ports to work. You can test for IP connectivity by pinging the IP address of the VMPS and verifying that you get a response.

Configuring Dynamic-Access Ports on VMPS Clients

If you are configuring a port on a cluster member switch as a dynamic-access port, first use the **rcommand** privileged EXEC command to log into the cluster member switch.



Caution

Dynamic-access port VLAN membership is for end stations or hubs connected to end stations. Connecting dynamic-access ports to other switches can cause a loss of connectivity.

Beginning in privileged EXEC mode, follow these steps to configure a dynamic-access port on a VMPS client switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the switch port that is connected to the end station.
Step 3	switchport mode access	Set the port to access mode.
Step 4	switchport access vlan dynamic	Configure the port as eligible for dynamic VLAN membership. The dynamic-access port must be connected to an end station.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries in the <i>Operational Mode</i> field of the display.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return an interface to its default configuration, use the **default interface** *interface-id* interface configuration command. To return an interface to its default switchport mode (dynamic auto), use the **no switchport mode** interface configuration command. To reset the access mode to the default VLAN for the switch, use the **no switchport access vlan** interface configuration command.

Reconfirming VLAN Memberships

Beginning in privileged EXEC mode, follow these steps to confirm the dynamic-access port VLAN membership assignments that the switch has received from the VMPS:

	Command	Purpose
Step 1	vmpls reconfirm	Reconfirm dynamic-access port VLAN membership.
Step 2	show vmpls	Verify the dynamic VLAN reconfirmation status.

Changing the Reconfirmation Interval

VMPS clients periodically reconfirm the VLAN membership information received from the VMPS. You can set the number of minutes after which reconfirmation occurs.

If you are configuring a member switch in a cluster, this parameter must be equal to or greater than the reconfirmation setting on the command switch. You must also first use the **rcommand** privileged EXEC command to log into the member switch.

Beginning in privileged EXEC mode, follow these steps to change the reconfirmation interval:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps reconfirm <i>minutes</i>	Enter the number of minutes between reconfirmations of the dynamic VLAN membership. The range is from 1 to 120. The default is 60 minutes.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmps	Verify the dynamic VLAN reconfirmation status in the <i>Reconfirm Interval</i> field of the display.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps reconfirm** global configuration command.

Changing the Retry Count

Beginning in privileged EXEC mode, follow these steps to change the number of times that the switch attempts to contact the VMPS before querying the next server:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vmps retry <i>count</i>	Change the retry count. The retry range is from 1 to 10; the default is 3.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vmps	Verify your entry in the <i>Server Retry Count</i> field of the display.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no vmps retry** global configuration command.

Monitoring the VMPS

You can display information about the VMPS by using the **show vmps** privileged EXEC command. The switch displays this information about the VMPS:

- VMPS VQP Version—the version of VQP used to communicate with the VMPS. The switch queries the VMPS that is using VQP version 1.
- Reconfirm Interval—the number of minutes the switch waits before reconfirming the VLAN-to-MAC-address assignments.
- Server Retry Count—the number of times VQP resends a query to the VMPS. If no response is received after this many tries, the switch starts to query the secondary VMPS.
- VMPS domain server—the IP address of the configured VLAN membership policy servers. The switch sends queries to the one marked *current*. The one marked *primary* is the primary server.
- VMPS Action—the result of the most recent reconfirmation attempt. A reconfirmation attempt can occur automatically when the reconfirmation interval expired, or you can force it by entering the **vmps reconfirm** privileged EXEC command or its CMS or SNMP equivalent

This is an example of output for the **show vmips** privileged EXEC command:

```
Switch# show vmips

VQP Client Status:
-----
VMPS VQP Version:    1
Reconfirm Interval:  60 min
Server Retry Count:  3
VMPS domain server:  172.20.128.86 (primary, current)
                    172.20.128.87

Reconfirmation status
-----
VMPS Action:         other
```

Troubleshooting Dynamic-Access Port VLAN Membership

The VMPS shuts down a dynamic-access port under these conditions:

- The VMPS is in secure mode, and it does not allow the host to connect to the port. The VMPS shuts down the port to prevent the host from connecting to the network.
- More than 20 active hosts reside on a dynamic-access port.

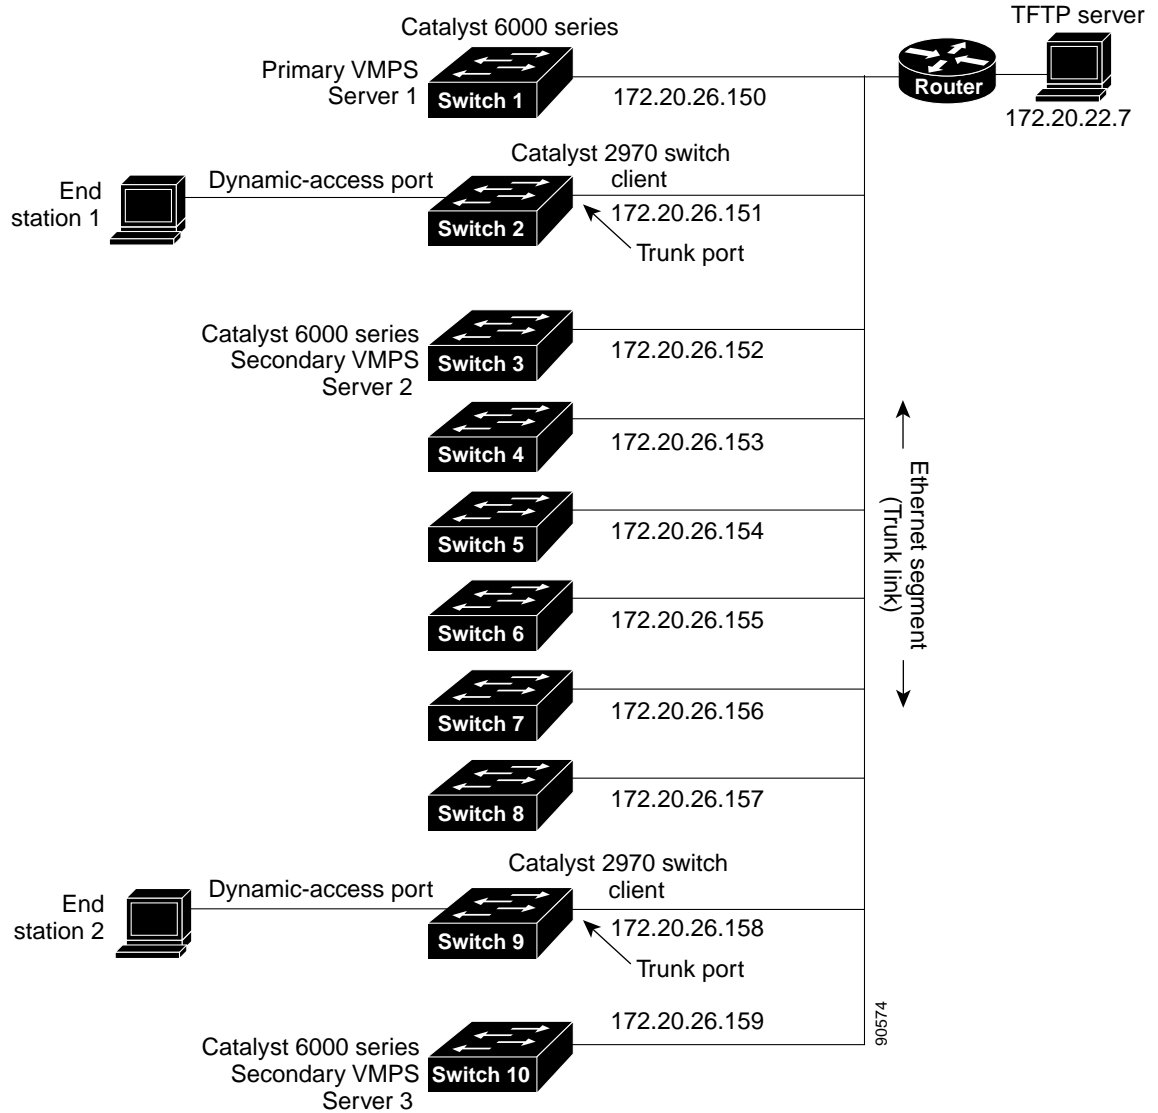
To re-enable a disabled dynamic-access port, enter the **shutdown** interface configuration command followed by the **no shutdown** interface configuration command.

VMPS Configuration Example

Figure 11-5 shows a network with a VMPS server switch and VMPS client switches with dynamic-access ports. In this example, these assumptions apply:

- The VMPS server and the VMPS client are separate switches.
- The Catalyst 6000 series Switch 1 is the primary VMPS server.
- The Catalyst 6000 series Switch 3 and Switch 10 are secondary VMPS servers.
- End stations are connected to the clients, Switch 2 and Switch 9.
- The database configuration file is stored on the TFTP server with the IP address 172.20.22.7.

Figure 11-5 Dynamic Port VLAN Membership Configuration





Configuring VTP

This chapter describes how to use the VLAN Trunking Protocol (VTP) and the VLAN database for managing VLANs with the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

The chapter includes these sections:

- [Understanding VTP, page 12-1](#)
- [Configuring VTP, page 12-6](#)
- [Monitoring VTP, page 12-15](#)

Understanding VTP

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches.

VTP is designed to work in an environment where updates are made on a single switch and are sent through VTP to other switches in the domain. It does not work well in a situation where multiple updates to the VLAN database occur simultaneously on switches in the same domain, which would result in an inconsistency in the VLAN database.

The switch supports 1005 VLANs, but the number of configured features affects the usage of the switch hardware. If the switch is notified by VTP of a new VLAN and the switch is already using the maximum available hardware resources, it sends a message that there are not enough hardware resources available and shuts down the VLAN. The output of the **show vlan** user EXEC command shows the VLAN in a suspended state.

VTP only learns about normal-range VLANs (VLAN IDs 1 to 1005). Extended-range VLANs (VLAN IDs greater than 1005) are not supported by VTP or stored in the VTP VLAN database.

This section contains information about these VTP parameters and characteristics.

- [The VTP Domain, page 12-2](#)
- [VTP Modes, page 12-3](#)
- [VTP Advertisements, page 12-3](#)
- [VTP Version 2, page 12-4](#)
- [VTP Pruning, page 12-4](#)

The VTP Domain

A VTP domain (also called a VLAN management domain) consists of one switch or several interconnected switches under the same administrative responsibility sharing the same VTP domain name. A switch can be in only one VTP domain. You make global VLAN configuration changes for the domain by using the command-line interface (CLI), Cluster Management Suite (CMS) software, or Simple Network Management Protocol (SNMP).

By default, the switch is in VTP no-management-domain state until it receives an advertisement for a domain over a trunk link (a link that carries the traffic of multiple VLANs) or until you configure a domain name. Until the management domain name is specified or learned, you cannot create or modify VLANs on a VTP server, and VLAN information is not propagated over the network.

If the switch receives a VTP advertisement over a trunk link, it inherits the management domain name and the VTP configuration revision number. The switch then ignores advertisements with a different domain name or an earlier configuration revision number.



Caution

Before adding a VTP client switch to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. If you add a switch that has a revision number higher than the revision number in the VTP domain, it can erase all VLAN information from the VTP server and VTP domain. See the [“Adding a VTP Client Switch to a VTP Domain” section on page 12-14](#) for the procedure for verifying and resetting the VTP configuration revision number.

When you make a change to the VLAN configuration on a VTP server, the change is propagated to all switches in the VTP domain. VTP advertisements are sent over all IEEE trunk connections, including Inter-Switch Link (ISL) and IEEE 802.1Q. VTP dynamically maps VLANs with unique names and internal index associates across multiple LAN types. Mapping eliminates excessive device administration required from network administrators.

If you configure a switch for VTP transparent mode, you can create and modify VLANs, but the changes are not sent to other switches in the domain, and they affect only the individual switch. However, configuration changes made when the switch is in this mode are saved in the switch running configuration and can be saved to the switch startup configuration file.

For domain name and password configuration guidelines, see the [“VTP Configuration Guidelines” section on page 12-7](#).

VTP Modes

You can configure a supported switch to be in one of the VTP modes listed in [Table 12-1](#).

Table 12-1 VTP Modes

VTP Mode	Description
VTP server	<p>In VTP server mode, you can create, modify, and delete VLANs, and specify other configuration parameters (such as the VTP version) for the entire VTP domain. VTP servers advertise their VLAN configurations to other switches in the same VTP domain and synchronize their VLAN configurations with other switches based on advertisements received over trunk links.</p> <p>In VTP server mode, VLAN configurations are saved in nonvolatile RAM (NVRAM). VTP server is the default mode.</p>
VTP client	<p>A VTP client behaves like a VTP server and transmits and receives VTP updates on its trunks, but you cannot create, change, or delete VLANs on a VTP client. VLANs are configured on another switch in the domain that is in server mode.</p> <p>In VTP client mode, VLAN configurations are not saved in NVRAM.</p>
VTP transparent	<p>VTP transparent switches do not participate in VTP. A VTP transparent switch does not advertise its VLAN configuration and does not synchronize its VLAN configuration based on received advertisements. However, in VTP version 2, transparent switches do forward VTP advertisements that they receive from other switches from their trunk interfaces. You can create, modify, and delete VLANs on a switch in VTP transparent mode. The switch must be in VTP transparent mode when you create extended-range VLANs. See the “Configuring Extended-Range VLANs” section on page 11-12.</p> <p>When the switch is in VTP transparent mode, the VTP and VLAN configurations are saved in NVRAM, but they are not advertised to other switches. In this mode, VTP mode and domain name are saved in the switch running configuration and you can save this information in the switch startup configuration file by entering the copy running-config startup-config privileged EXEC command.</p>

VTP Advertisements

Each switch in the VTP domain sends periodic global configuration advertisements from each trunk port to a reserved multicast address. Neighboring switches receive these advertisements and update their VTP and VLAN configurations as necessary.



Note

Because trunk ports send and receive VTP advertisements, you must ensure that at least one trunk port is configured on the switch and that this trunk port is connected to the trunk port of another switch. Otherwise, the switch cannot receive any VTP advertisements. For more information on trunk ports, see the [“Configuring VLAN Trunks”](#) section on page 11-15.

VTP advertisements distribute this global domain information:

- VTP domain name
- VTP configuration revision number
- Update identity and update timestamp
- MD5 digest VLAN configuration, including maximum transmission unit (MTU) size for each VLAN.
- Frame format

VTP advertisements distribute this VLAN information for each configured VLAN:

- VLAN IDs (ISL and 802.1Q)
- VLAN name
- VLAN type
- VLAN state
- Additional VLAN configuration information specific to the VLAN type

VTP Version 2

If you use VTP in your network, you must decide whether to use version 1 or version 2. By default, VTP operates in version 1.

VTP version 2 supports these features not supported in version 1:

- Token Ring support—VTP version 2 supports Token Ring Bridge Relay Function (TrBRF) and Token Ring Concentrator Relay Function (TrCRF) VLANs. For more information about Token Ring VLANs, see the [“Configuring Normal-Range VLANs” section on page 11-4](#).
- Unrecognized Type-Length-Value (TLV) support—A VTP server or client propagates configuration changes to its other trunks, even for TLVs it is not able to parse. The unrecognized TLV is saved in NVRAM when the switch is operating in VTP server mode.
- Version-Dependent Transparent Mode—In VTP version 1, a VTP transparent switch inspects VTP messages for the domain name and version and forwards a message only if the version and domain name match. Because VTP version 2 supports only one domain, it forwards VTP messages in transparent mode without inspecting the version and domain name.
- Consistency Checks—In VTP version 2, VLAN consistency checks (such as VLAN names and values) are performed only when you enter new information through the CLI, the Cluster Management Software (CMS), or SNMP. Consistency checks are not performed when new information is obtained from a VTP message or when information is read from NVRAM. If the MD5 digest on a received VTP message is correct, its information is accepted.

VTP Pruning

VTP pruning increases network available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to reach the destination devices. Without VTP pruning, a switch floods broadcast, multicast, and unknown unicast traffic across all trunk links within a VTP domain even though receiving switches might discard them. VTP pruning is disabled by default.

VTP pruning blocks unneeded flooded traffic to VLANs on trunk ports that are included in the pruning-eligible list. Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible switch trunk ports. If the VLANs are configured as pruning-ineligible, the flooding continues. VTP pruning is supported with VTP version 1 and version 2.

[Figure 12-1](#) shows a switched network without VTP pruning enabled. Port 1 on Switch A and Port 2 on Switch D are assigned to the Red VLAN. If a broadcast is sent from the host connected to Switch A, Switch A floods the broadcast and every switch in the network receives it, even though Switches C, E, and F have no ports in the Red VLAN.

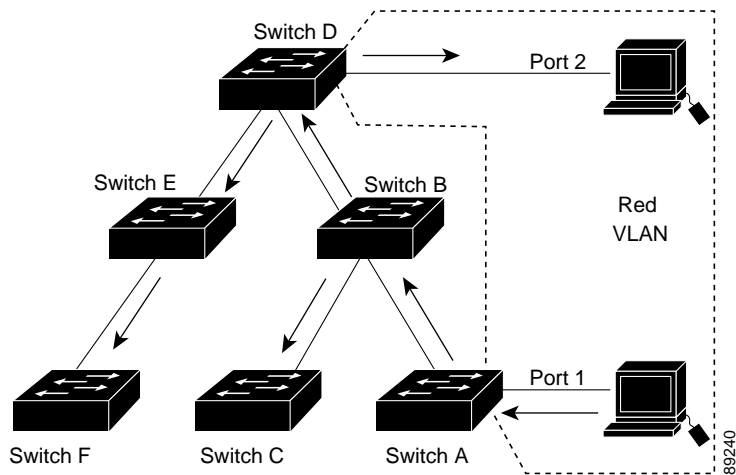
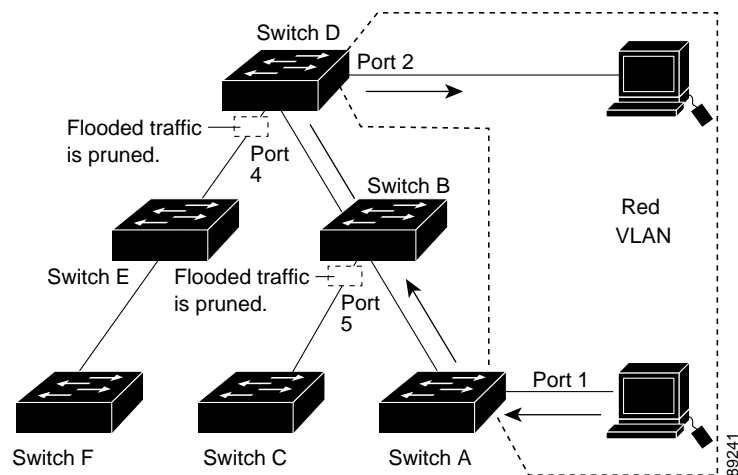
Figure 12-1 Flooding Traffic without VTP Pruning

Figure 12-2 shows a switched network with VTP pruning enabled. The broadcast traffic from Switch A is not forwarded to Switches C, E, and F because traffic for the Red VLAN has been pruned on the links shown (Port 5 on Switch B and Port 4 on Switch D).

Figure 12-2 Optimized Flooded Traffic with VTP Pruning

Enabling VTP pruning on a VTP server enables pruning for the entire management domain. Making VLANs pruning-eligible or pruning-ineligible affects pruning eligibility for those VLANs on that trunk only (not on all switches in the VTP domain).

See the “[Enabling VTP Pruning](#)” section on page 12-13. VTP pruning takes effect several seconds after you enable it. VTP pruning does not prune traffic from VLANs that are pruning-ineligible. VLAN 1 and VLANs 1002 to 1005 are always pruning-ineligible; traffic from these VLANs cannot be pruned. Extended-range VLANs (VLAN IDs higher than 1005) are also pruning-ineligible.

VTP pruning is not designed to function in VTP transparent mode. If one or more switches in the network are in VTP transparent mode, you should do one of these:

- Turn off VTP pruning in the entire network.

- Turn off VTP pruning by making all VLANs on the trunk of the switch upstream to the VTP transparent switch pruning ineligible.

To configure VTP pruning on an interface, use the **switchport trunk pruning vlan** interface configuration command (see the “[Changing the Pruning-Eligible List](#)” section on page 11-21). VTP pruning operates when an interface is trunking. You can set VLAN pruning-eligibility, whether or not VTP pruning is enabled for the VTP domain, whether or not any given VLAN exists, and whether or not the interface is currently trunking.

Configuring VTP

This section includes guidelines and procedures for configuring VTP. These sections are included:

- [Default VTP Configuration, page 12-6](#)
- [VTP Configuration Options, page 12-6](#)
- [VTP Configuration Guidelines, page 12-7](#)
- [Configuring a VTP Server, page 12-9](#)
- [Configuring a VTP Client, page 12-10](#)
- [Disabling VTP \(VTP Transparent Mode\), page 12-11](#)
- [Enabling VTP Version 2, page 12-12](#)
- [Enabling VTP Pruning, page 12-13](#)
- [Adding a VTP Client Switch to a VTP Domain, page 12-14](#)

Default VTP Configuration

[Table 12-2](#) shows the default VTP configuration.

Table 12-2 *Default VTP Configuration*

Feature	Default Setting
VTP domain name	Null.
VTP mode	Server.
VTP version	Version 1 (version 2 is disabled).
VTP password	None.
VTP pruning	Disabled.

VTP Configuration Options

You can configure VTP by using these configuration modes.

- [VTP Configuration in Global Configuration Mode, page 12-7](#)
- [VTP Configuration in VLAN Database Configuration Mode, page 12-7](#)

You access VLAN database configuration mode by entering the **vlan database** privileged EXEC command.

For detailed information about **vtp** commands, refer to the command reference for this release.

VTP Configuration in Global Configuration Mode

You can use the **vtp** global configuration command to set the VTP password, the version, the VTP file name, the interface providing updated VTP information, the domain name, and the mode, and to disable or enable pruning. For more information about available keywords, refer to the command descriptions in the command reference for this release. The VTP information is saved in the VTP VLAN database. When VTP mode is transparent, the VTP domain name and mode are also saved in the switch running configuration file, and you can save it in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command. You must use this command if you want to save VTP mode as transparent, even if the switch resets.

When you save VTP information in the switch startup configuration file and reboot the switch, the switch configuration is determined as follows:

- If the VTP mode is transparent in the startup configuration and the VLAN database and the VTP domain name from the VLAN database matches that in the startup configuration file, the VLAN database is ignored (cleared), and the VTP and VLAN configurations in the startup configuration file are used. The VLAN database revision number remains unchanged in the VLAN database.
- If the VTP mode or domain name in the startup configuration do not match the VLAN database, the domain name and VTP mode and configuration for the first 1005 VLANs use the VLAN database information.

VTP Configuration in VLAN Database Configuration Mode

You can configure all VTP parameters in VLAN database configuration mode, which you access by entering the **vlan database** privileged EXEC command. For more information about available keywords, refer to the **vtp** VLAN database configuration command description in the command reference for this release. When you enter the **exit** command in VLAN database configuration mode, it applies all the commands that you entered and updates the VLAN database. VTP messages are sent to other switches in the VTP domain, and the privileged EXEC mode prompt appears.

If VTP mode is transparent, the domain name and the mode (transparent) are saved in the switch running configuration, and you can save this information in the switch startup configuration file by entering the **copy running-config startup-config** privileged EXEC command.

VTP Configuration Guidelines

These sections describe guidelines you should follow when implementing VTP in your network.

Domain Names

When configuring VTP for the first time, you must always assign a domain name. You must configure all switches in the VTP domain with the same domain name. Switches in VTP transparent mode do not exchange VTP messages with other switches, and you do not need to configure a VTP domain name for them.



Note

If NVRAM and DRAM storage is sufficient, all switches in a VTP domain should be in VTP server mode.

**Caution**

Do not configure a VTP domain if all switches are operating in VTP client mode. If you configure the domain, it is impossible to make changes to the VLAN configuration of that domain. Make sure that you configure at least one switch in the VTP domain for VTP server mode.

Passwords

You can configure a password for the VTP domain, but it is not required. If you do configure a domain password, all domain switches must share the same password and you must configure the password on each switch in the management domain. Switches without a password or with the wrong password reject VTP advertisements.

If you configure a VTP password for a domain, a switch that is booted without a VTP configuration does not accept VTP advertisements until you configure it with the correct password. After the configuration, the switch accepts the next VTP advertisement that uses the same password and domain name in the advertisement.

If you are adding a new switch to an existing network with VTP capability, the new switch learns the domain name only after the applicable password has been configured on it.

**Caution**

When you configure a VTP domain password, the management domain does not function properly if you do not assign a management domain password to each switch in the domain.

VTP Version

Follow these guidelines when deciding which VTP version to implement:

- All switches in a VTP domain must run the same VTP version.
- A VTP version 2-capable switch can operate in the same VTP domain as a switch running VTP version 1 if version 2 is disabled on the version 2-capable switch (version 2 is disabled by default).
- Do not enable VTP version 2 on a switch unless all of the switches in the same VTP domain are version-2-capable. When you enable version 2 on a switch, all of the version-2-capable switches in the domain enable version 2. If there is a version 1-only switch, it does not exchange VTP information with switches with version 2 enabled.
- If there are TrBRF and TrCRF Token Ring networks in your environment, you must enable VTP version 2 for Token Ring VLAN switching to function properly. To run Token Ring and Token Ring-Net, disable VTP version 2.

Configuration Requirements

When you configure VTP, you must configure a trunk port so that the switch can send and receive VTP advertisements to and from other switches in the domain.

For more information, see the [“Configuring VLAN Trunks” section on page 11-15](#).

If you are configuring VTP on a cluster member switch to a VLAN, use the **rcommand** privileged EXEC command to log into the member switch. For more information about the command, refer to the command reference for this release.

If you are configuring extended-range VLANs on the switch, the switch must be in VTP transparent mode.

Configuring a VTP Server

When a switch is in VTP server mode, you can change the VLAN configuration and have it propagated throughout the network.



Note

If extended-range VLANs are configured on the switch, you cannot change VTP mode to server. You receive an error message, and the configuration is not allowed.

Beginning in privileged EXEC mode, follow these steps to configure the switch as a VTP server:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode server	Configure the switch for VTP server mode (the default).
Step 3	vtp domain <i>domain-name</i>	Configure the VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	vtp password <i>password</i>	(Optional) Set the password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
Step 5	end	Return to privileged EXEC mode.
Step 6	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

To return the switch to a no-password state, use the **no vtp password** global configuration command.

This example shows how to use global configuration mode to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch# config terminal
Switch(config)# vtp mode server
Switch(config)# vtp domain eng_group
Switch(config)# vtp password mypassword
Switch(config)# end
```

You can also use VLAN database configuration mode to configure VTP parameters.

Beginning in privileged EXEC mode, follow these steps to use VLAN database configuration mode to configure the switch as a VTP server:

	Command	Purpose
Step 1	vlan database	Enter VLAN database configuration mode.
Step 2	vtp server	Configure the switch for VTP server mode (the default).

	Command	Purpose
Step 3	vtp domain <i>domain-name</i>	Configure a VTP administrative-domain name. The name can be from 1 to 32 characters. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	vtp password <i>password</i>	(Optional) Set a password for the VTP domain. The password can be from 8 to 64 characters. If you configure a VTP password, the VTP domain does not function properly if you do not assign the same password to each switch in the domain.
Step 5	exit	Update the VLAN database, propagate it throughout the administrative domain, and return to privileged EXEC mode.
Step 6	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

To return the switch to a no-password state, use the **no vtp password** VLAN database configuration command.

This example shows how to use VLAN database configuration mode to configure the switch as a VTP server with the domain name *eng_group* and the password *mypassword*:

```
Switch# vlan database
Switch(vlan)# vtp server
Switch(vlan)# vtp domain eng_group
Switch(vlan)# vtp password mypassword
Switch(vlan)# exit
APPLY completed.
Exiting...
Switch#
```

Configuring a VTP Client

When a switch is in VTP client mode, you cannot change its VLAN configuration. The client switch receives VTP updates from a VTP server in the VTP domain and then modifies its configuration accordingly.



Note

If extended-range VLANs are configured on the switch, you cannot change VTP mode to client. You receive an error message, and the configuration is not allowed.



Caution

If all switches are operating in VTP client mode, do not configure a VTP domain name. If you do, it is impossible to make changes to the VLAN configuration of that domain. Therefore, make sure you configure at least one switch as a VTP server.

Beginning in privileged EXEC mode, follow these steps to configure the switch as a VTP client:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode client	Configure the switch for VTP client mode. The default setting is VTP server.
Step 3	vtp domain <i>domain-name</i>	(Optional) Enter the VTP administrative-domain name. The name can be from 1 to 32 characters. This should be the same domain name as the VTP server. All switches operating in VTP server or client mode under the same administrative responsibility must be configured with the same domain name.
Step 4	vtp password <i>password</i>	(Optional) Enter the password for the VTP domain.
Step 5	end	Return to privileged EXEC mode.
Step 6	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.

Use the **no vtp mode** global configuration command to return the switch to VTP server mode. To return the switch to a no-password state, use the **no vtp password** privileged EXEC command. When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.



Note

You can also configure a VTP client by using the **vlan database** privileged EXEC command to enter VLAN database configuration mode and entering the **vtp client** command, similar to the second procedure under “[Configuring a VTP Server](#)” section on page 12-9. Use the **no vtp client** VLAN database configuration command to return the switch to VTP server mode or the **no vtp password** VLAN database configuration command to return the switch to a no-password state. When you configure a domain name, it cannot be removed; you can only reassign a switch to a different domain.

Disabling VTP (VTP Transparent Mode)

When you configure the switch for VTP transparent mode, you disable VTP on the switch. The switch does not send VTP updates and does not act on VTP updates received from other switches. However, a VTP transparent switch running VTP version 2 does forward received VTP advertisements on all of its trunk links.



Note

Before you create extended-range VLANs (VLAN IDs 1006 to 4094), you must set VTP mode to transparent by using the **vtp mode transparent** global configuration command. Save this configuration to the startup configuration so that the switch boots up in VTP transparent mode. Otherwise, you lose the extended-range VLAN configuration if the switch resets and boots up in VTP server mode (the default).

Beginning in privileged EXEC mode, follow these steps to configure VTP transparent mode and save the VTP configuration in the switch startup configuration file:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp mode transparent	Configure the switch for VTP transparent mode (disable VTP).
Step 3	end	Return to privileged EXEC mode.
Step 4	show vtp status	Verify your entries in the <i>VTP Operating Mode</i> and the <i>VTP Domain Name</i> fields of the display.
Step 5	copy running-config startup-config	(Optional) Save the configuration in the startup configuration file. Note Only VTP mode and domain name are saved in the switch running configuration and can be copied to the startup configuration file.

To return the switch to VTP server mode, use the **no vtp mode** global configuration command.



Note

If extended-range VLANs are configured on the switch, you cannot change the VTP mode to server. You receive an error message, and the configuration is not allowed.



Note

You can also configure VTP transparent mode by using the **vlan database** privileged EXEC command to enter VLAN database configuration mode and by entering the **vtp transparent** command, similar to the second procedure under the “[Configuring a VTP Server](#)” section on [page 12-9](#). Use the **no vtp transparent** VLAN database configuration command to return the switch to VTP server mode. If extended-range VLANs are configured on the switch, you cannot change VTP mode to server. You receive an error message, and the configuration is not allowed.

Enabling VTP Version 2

VTP version 2 is disabled by default on VTP version 2-capable switches. When you enable VTP version 2 on a switch, every VTP version 2-capable switch in the VTP domain enables version 2. You can only configure the version on switches in VTP server or transparent mode.



Caution

VTP version 1 and VTP version 2 are not interoperable on switches in the same VTP domain. Every switch in the VTP domain must use the same VTP version. Do not enable VTP version 2 unless every switch in the VTP domain supports version 2.



Note

In TrCRF and TrBRF Token ring environments, you must enable VTP version 2 for Token Ring VLAN switching to function properly. For Token Ring and Token Ring-Net media, VTP version 2 must be disabled.

For more information on VTP version configuration guidelines, see the “[VTP Version](#)” section on [page 12-8](#).

Beginning in privileged EXEC mode, follow these steps to enable VTP version 2:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp version 2	Enable VTP version 2 on the switch. VTP version 2 is disabled by default on VTP version 2-capable switches.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vtp status	Verify that VTP version 2 is enabled in the <i>VTP V2 Mode</i> field of the display.

To disable VTP version 2, use the **no vtp version** global configuration command.



Note

You can also enable VTP version 2 by using the **vlan database** privileged EXEC command to enter VLAN database configuration mode and entering the **vtp v2-mode** VLAN database configuration command. To disable VTP version 2, use the **no vtp v2-mode** VLAN database configuration command.

Enabling VTP Pruning

Pruning increases available bandwidth by restricting flooded traffic to those trunk links that the traffic must use to access the destination devices. You can only enable VTP pruning on a switch in VTP server mode.

Beginning in privileged EXEC mode, follow these steps to enable VTP pruning in the VTP domain:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vtp pruning	Enable pruning in the VTP administrative domain. By default, pruning is disabled. You need to enable pruning on only one switch in VTP server mode.
Step 3	end	Return to privileged EXEC mode.
Step 4	show vtp status	Verify your entries in the <i>VTP Pruning Mode</i> field of the display.

To disable VTP pruning, use the **no vtp pruning** global configuration command.



Note

You can also enable VTP pruning by using the **vlan database** privileged EXEC command to enter VLAN database configuration mode and entering the **vtp pruning** VLAN database configuration command. To disable VTP pruning, use the **no vtp pruning** VLAN database configuration command. You can also enable VTP version 2 by using the **vtp pruning** privileged EXEC command. However, this command will not be available in future releases.

Pruning is supported with VTP version 1 and version 2. If you enable pruning on the VTP server, it is enabled for the entire VTP domain.

Only VLANs included in the pruning-eligible list can be pruned. By default, VLANs 2 through 1001 are pruning eligible on trunk ports. Reserved VLANs and extended-range VLANs cannot be pruned. To change the pruning-eligible VLANs, see the [“Changing the Pruning-Eligible List”](#) section on page 11-21.

Adding a VTP Client Switch to a VTP Domain

Before adding a VTP client to a VTP domain, always verify that its VTP configuration revision number is *lower* than the configuration revision number of the other switches in the VTP domain. Switches in a VTP domain always use the VLAN configuration of the switch with the highest VTP configuration revision number. If you add a switch that has a revision number higher than the revision number in the VTP domain, it can erase all VLAN information from the VTP server and VTP domain.

Beginning in privileged EXEC mode, follow these steps to verify and reset the VTP configuration revision number on a switch *before* adding it to a VTP domain:

	Command	Purpose
Step 1	show vtp status	Check the VTP configuration revision number. If the number is 0, add the switch to the VTP domain. If the number is greater than 0, follow these steps: a. Write down the domain name. b. Write down the configuration revision number. c. Continue with the next steps to reset the switch configuration revision number.
Step 2	configure terminal	Enter global configuration mode.
Step 3	vtp domain <i>domain-name</i>	Change the domain name from the original one displayed in Step 1 to a new name.
Step 4	end	The VLAN information on the switch is updated and the configuration revision number is reset to 0. You return to privileged EXEC mode.
Step 5	show vtp status	Verify that the configuration revision number has been reset to 0.
Step 6	configure terminal	Enter global configuration mode.
Step 7	vtp domain <i>domain-name</i>	Enter the original domain name on the switch.
Step 8	end	The VLAN information on the switch is updated, and you return to privileged EXEC mode.
Step 9	show vtp status	(Optional) Verify that the domain name is the same as in Step 1 and that the configuration revision number is 0.

You can also change the VTP domain name by entering the **vlan database** privileged EXEC command to enter VLAN database configuration mode and by entering the **vtp domain** *domain-name* command. In this mode, you must enter the **exit** command to update VLAN information and return to privileged EXEC mode.

After resetting the configuration revision number, add the switch to the VTP domain.



Note

You can use the **vtp mode transparent** global configuration command or the **vtp transparent** VLAN database configuration command to disable VTP on the switch, and then change its VLAN information without affecting the other switches in the VTP domain.

Monitoring VTP

You monitor VTP by displaying VTP configuration information: the domain name, the current VTP revision, and the number of VLANs. You can also display statistics about the advertisements sent and received by the switch.

Table 12-3 shows the privileged EXEC commands for monitoring VTP activity.

Table 12-3 VTP Monitoring Commands

Command	Purpose
show vtp status	Display the VTP switch configuration information.
show vtp counters	Display counters about VTP messages that have been sent and received.



Configuring Voice VLAN

This chapter describes how to configure the voice VLAN feature on the Catalyst 2970 switch. Voice VLAN is referred to as an *auxiliary VLAN* in some Catalyst 6000 family switch documentation.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding Voice VLAN, page 13-1](#)
- [Configuring Voice VLAN, page 13-3](#)
- [Displaying Voice VLAN, page 13-6](#)

Understanding Voice VLAN

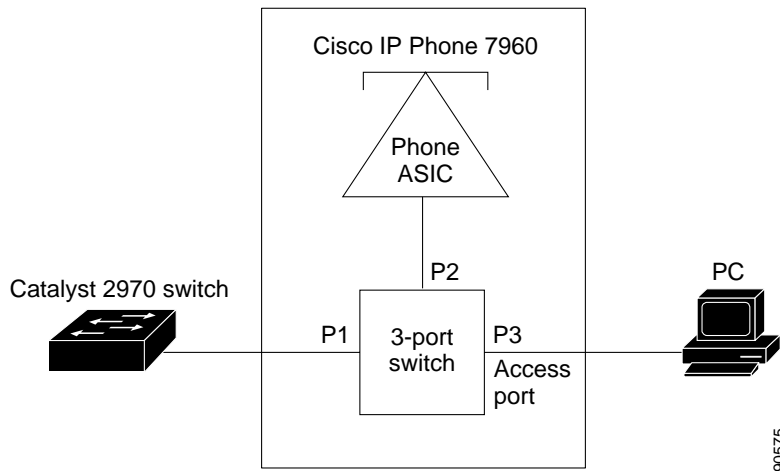
The voice VLAN feature enables access ports to carry IP voice traffic from an IP phone. When the switch is connected to a Cisco 7960 IP Phone, the IP Phone sends voice traffic with Layer 3 IP precedence and Layer 2 class of service (CoS) values, which are both set to 5 by default. Because the sound quality of an IP phone call can deteriorate if the data is unevenly sent, the switch supports quality of service (QoS) based on IEEE 802.1P CoS. QoS uses classification and scheduling to send network traffic from the switch in a predictable manner. For more information on QoS, see [Chapter 27, “Configuring QoS.”](#)

The Cisco 7960 IP Phone is a configurable device, and you can configure it to forward traffic with an 802.1P priority. You can configure the switch to trust or override the traffic priority assigned by an IP Phone.

The Cisco IP Phone contains an integrated three-port 10/100 switch as shown in [Figure 13-1](#). The ports provide dedicated connections to these devices:

- Port 1 connects to the switch or other voice-over-IP (VoIP) device.
- Port 2 is an internal 10/100 interface that carries the IP phone traffic.
- Port 3 (access port) connects to a PC or other device.

[Figure 13-1](#) shows one way to connect a Cisco 7960 IP Phone.

Figure 13-1 Cisco 7960 IP Phone Connected to a Switch

Cisco IP Phone Voice Traffic

You can configure an access port with an attached Cisco IP Phone to use one VLAN for voice traffic and another VLAN for data traffic from a device attached to the phone. You can configure access ports on the switch to send Cisco Discovery Protocol (CDP) packets that instruct an attached Cisco IP Phone to send voice traffic to the switch in any of these ways:

- In the voice VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN tagged with a Layer 2 CoS priority value
- In the access VLAN, untagged (no Layer 2 CoS priority value)



Note

In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5 for voice traffic and 3 for voice control traffic).

Cisco IP Phone Data Traffic

The switch can also process tagged data traffic (traffic in 802.1Q or 802.1P frame types) from the device attached to the access port on the Cisco IP Phone (see [Figure 13-1](#)). You can configure Layer 2 access ports on the switch to send CDP packets that instruct the attached Cisco IP Phone to configure the IP phone access port in one of these modes:

- In trusted mode, all traffic received through the access port on the Cisco IP Phone passes through the IP phone unchanged.
- In untrusted mode, all traffic in 802.1Q or 802.1P frames received through the access port on the IP phone receive a configured Layer 2 CoS value. The default Layer 2 CoS value is 0. Untrusted mode is the default.



Note

Untagged traffic from the device attached to the Cisco IP Phone passes through the IP phone unchanged, regardless of the trust state of the access port on the IP phone.

Configuring Voice VLAN

This section describes how to configure voice VLAN on access ports. This section contains this configuration information:

- [Default Voice VLAN Configuration, page 13-3](#)
- [Voice VLAN Configuration Guidelines, page 13-3](#)
- [Configuring a Port Connected to a Cisco 7960 IP Phone, page 13-4](#)

Default Voice VLAN Configuration

The voice VLAN feature is disabled by default.

When the voice VLAN feature is enabled, all untagged traffic is sent according to the default CoS priority of the port.

The CoS value is not trusted for 802.1P or 802.1Q tagged traffic.

Voice VLAN Configuration Guidelines

These are the voice VLAN configuration guidelines:



Note

You should configure voice VLAN on switch access ports; voice VLAN is not supported on trunk ports. Voice VLAN is only supported on access ports and not on trunk ports, even though the configuration is allowed.

- Before you enable voice VLAN, we recommend that you enable QoS on the switch by entering the **mls qos** global configuration command and configure the port trust state to trust by entering the **mls qos trust cos** interface configuration command. If you use the auto-QoS feature, these settings are automatically configured. For more information, see [Chapter 27, “Configuring QoS.”](#)
- You must enable CDP on the switch port connected to the Cisco IP Phone to send configuration to the Cisco IP Phone. (CDP is enabled by default globally and on all switch interfaces.)
- The Port Fast feature is automatically enabled when voice VLAN is configured. When you disable voice VLAN, the Port Fast feature is not automatically disabled.
- If the Cisco IP Phone and a device attached to the Cisco IP Phone are in the same VLAN, they must be in the same IP subnet. These conditions indicate that they are in the same VLAN:
 - They both use 802.1p or untagged frames.
 - The Cisco IP Phone uses 802.1p frames and the device uses untagged frames.
 - The Cisco IP Phone uses untagged frames and the device uses 802.1p frames.
 - The Cisco IP Phone uses 802.1Q frames and the voice VLAN is the same as the access VLAN.
- The Cisco IP Phone and a device attached to the phone cannot communicate if they are in the same VLAN and subnet but use different frame types because traffic in the same subnet is not routed (routing would eliminate the frame type difference).
- You cannot configure static secure MAC addresses in the voice VLAN.

- Voice VLAN ports can also be these port types:
 - Dynamic access port. See the “[Configuring Dynamic-Access Ports on VMPS Clients](#)” section on page 11-29 for more information.
 - 802.1X authenticated port. See the “[Configuring 802.1X Authentication](#)” section on page 8-12 for more information.
 - Protected port. See the “[Configuring Protected Ports](#)” section on page 19-5 for more information.
 - A source or destination port for a SPAN or RSPAN session.
 - Secure port. See the “[Configuring Port Security](#)” section on page 19-7 for more information.

**Note**

When you enable port security on an interface that is also configured with a voice VLAN, you must set the maximum allowed secure addresses on the port to two plus the maximum number of secure addresses allowed on the access VLAN. When the port is connected to a Cisco IP phone, the IP phone requires up to two MAC addresses. The IP phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the IP phone requires additional MAC addresses.

Configuring a Port Connected to a Cisco 7960 IP Phone

Because a Cisco 7960 IP Phone also supports a connection to a PC or other device, a port connecting the switch to a Cisco IP Phone can carry mixed traffic. You can configure a port to determine how the IP phone carries voice traffic and data traffic.

This section includes these topics:

- [Configuring IP Phone Voice Traffic](#), page 13-4
- [Configuring the Priority of Incoming Data Frames](#), page 13-5

Configuring IP Phone Voice Traffic

You can configure a port connected to the Cisco IP Phone to send CDP packets to the phone to configure the way in which the phone sends voice traffic. The phone can carry voice traffic in 802.1Q frames for a specified voice VLAN with a Layer 2 CoS value. It can use 802.1P priority tagging to give voice traffic a higher priority and forward all voice traffic through the native (access) VLAN. The IP phone can also send untagged voice traffic or use its own configuration to send voice traffic in the access VLAN. In all configurations, the voice traffic carries a Layer 3 IP precedence value (the default is 5).

Beginning in privileged EXEC mode, follow these steps to configure voice traffic on a port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the IP phone.
Step 3	mls qos trust cos	Configure the interface to classify ingress traffic packets by using the packet CoS value. For untagged packets, the port default CoS value is used.
		Note Before configuring the port trust state, you must first globally enable QoS by using the mls qos global configuration command.

	Command	Purpose
Step 4	switchport voice vlan { <i>vlan-id</i> / dot1p / none / untagged }	Configure how the Cisco IP Phone carries voice traffic: <ul style="list-style-type: none"> vlan-id—Configure the Cisco IP Phone to forward all voice traffic through the specified VLAN. By default, the Cisco IP Phone forwards the voice traffic with an 802.1Q priority of 5. Valid VLAN IDs are from 1 to 4094. dot1p—Configure the Cisco IP Phone to use 802.1P priority tagging for voice traffic and to use the default native VLAN (VLAN 0) to carry all traffic. By default, the Cisco IP Phone forwards the voice traffic with an 802.1P priority of 5. none—Allow the IP phone to use its own configuration to send untagged voice traffic. untagged—Configure the phone to send untagged voice traffic.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport or show running-config interface <i>interface-id</i>	Verify your voice VLAN entries. Verify your QoS and voice VLAN entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the port to its default setting, use the **no switchport voice vlan** interface configuration command.

Configuring the Priority of Incoming Data Frames

You can connect a PC or other data device to a Cisco IP Phone port. To process tagged data traffic (in 802.1Q or 802.1P frames), you can configure the switch to send CDP packets to instruct the IP phone how to send data packets from the device attached to the access port on the Cisco IP Phone. The PC can generate packets with an assigned CoS value. You can configure the Cisco IP Phone to not change (trust) or to override (not trust) the priority of frames arriving on the IP phone port from connected devices.

Beginning in privileged EXEC mode, follow these steps to set the priority of data traffic received from the nonvoice port on the Cisco IP Phone:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the IP phone.
Step 3	switchport priority extend { <i>cos value</i> trust }	Set the priority of data traffic received from the IP phone access port: <ul style="list-style-type: none"> cos value—Configure the IP phone to override the priority received from the PC or the attached device with the specified CoS value. The value is a number from 0 to 7, with 7 as the highest priority. The default priority is cos 0. trust—Configure the IP phone access port to trust the priority received from the PC or the attached device.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the port to its default setting, use the **no switchport priority extend** interface configuration command.

Displaying Voice VLAN

To display voice VLAN configuration for an interface, use the **show interfaces** *interface-id* **switchport** privileged EXEC command.



Configuring STP

This chapter describes how to configure the Spanning Tree Protocol (STP) on port-based VLANs on the Catalyst 2970 switch. The switch uses the per-VLAN spanning-tree plus (PVST+) protocol based on the IEEE 802.1D standard and Cisco proprietary extensions, or it can use the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol based on the IEEE 802.1W standard.

For information about the Multiple Spanning Tree Protocol (MSTP) and how to map multiple VLANs to the same spanning-tree instance, see [Chapter 15, “Configuring MSTP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 16, “Configuring Optional Spanning-Tree Features.”](#)



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding Spanning-Tree Features, page 14-1](#)
- [Configuring Spanning-Tree Features, page 14-11](#)
- [Displaying the Spanning-Tree Status, page 14-22](#)

Understanding Spanning-Tree Features

These sections describe how basic spanning-tree features work:

- [STP Overview, page 14-2](#)
- [Spanning-Tree Topology and BPDUs, page 14-3](#)
- [Bridge ID, Switch Priority, and Extended System ID, page 14-4](#)
- [Spanning-Tree Interface States, page 14-4](#)
- [How a Switch or Port Becomes the Root Switch or Root Port, page 14-7](#)
- [Spanning Tree and Redundant Connectivity, page 14-8](#)
- [Spanning-Tree Address Management, page 14-8](#)
- [Accelerated Aging to Retain Connectivity, page 14-8](#)
- [Spanning-Tree Modes and Protocols, page 14-9](#)
- [Supported Spanning-Tree Instances, page 14-9](#)

- [Spanning-Tree Interoperability and Backward Compatibility, page 14-10](#)
- [STP and IEEE 802.1Q Trunks, page 14-10](#)

For configuration information, see the [“Configuring Spanning-Tree Features” section on page 14-11](#).

For information about optional spanning-tree features, see [Chapter 16, “Configuring Optional Spanning-Tree Features.”](#)

STP Overview

STP is a Layer 2 link management protocol that provides path redundancy while preventing loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Switches might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an unstable network. Spanning-tree operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

The STP uses a spanning-tree algorithm to select one switch of a redundantly connected network as the root of the spanning tree. The algorithm calculates the best loop-free path through a switched Layer 2 network by assigning a role to each port based on the role of the port in the active topology:

- Root—A forwarding port elected for the spanning-tree topology
- Designated—A forwarding port elected for every switched LAN segment
- Alternate—A blocked port providing an alternate path to the root port in the spanning tree
- Backup—A blocked port in a loopback configuration

Switches that have ports with these assigned roles are called root or designated switches.

Spanning tree forces redundant data paths into a standby (blocked) state. If a network segment in the spanning tree fails and a redundant path exists, the spanning-tree algorithm recalculates the spanning-tree topology and activates the standby path. Switches send and receive spanning-tree frames, called bridge protocol data units (BPDUs), at regular intervals. The switches do not forward these frames but use them to construct a loop-free path. BPDUs contain information about the sending switch and its ports, including switch and MAC addresses, switch priority, port priority, and path cost. Spanning tree uses this information to elect the root switch and root port for the switched network and the root port and designated port for each switched segment.

When two interfaces on a switch are part of a loop, the spanning-tree port priority and path cost settings determine which interface is put in the forwarding state and which is put in the blocking state. The spanning-tree port priority value represents the location of an interface in the network topology and how well it is located to pass traffic. The path cost value represents the media speed.

Spanning-Tree Topology and BPDUs

The stable, active spanning-tree topology of a switched network is determined by these elements:

- The unique bridge ID (switch priority and MAC address) associated with each VLAN on each switch.
- The spanning-tree path cost to the root switch.
- The port identifier (port priority and MAC address) associated with each Layer 2 interface.

When the switches in a network are powered up, each functions as the root switch. Each switch sends a configuration BPDU through all of its ports. The BPDUs communicate and compute the spanning-tree topology. Each configuration BPDU contains this information:

- The unique bridge ID of the switch that the sending switch identifies as the root switch
- The spanning-tree path cost to the root
- The bridge ID of the sending switch
- Message age
- The identifier of the sending interface
- Values for the hello, forward delay, and max-age protocol timers

When a switch receives a configuration BPDU that contains *superior* information (lower bridge ID, lower path cost, and so forth), it stores the information for that port. If this BPDU is received on the root port of the switch, the switch also forwards it with an updated message to all attached LANs for which it is the designated switch.

If a switch receives a configuration BPDU that contains *inferior* information to that currently stored for that port, it discards the BPDU. If the switch is a designated switch for the LAN from which the inferior BPDU was received, it sends that LAN a BPDU containing the up-to-date information stored for that port. In this way, inferior information is discarded, and superior information is propagated on the network.

A BPDU exchange results in these actions:

One switch in the network is elected as the root switch (the logical center of the spanning-tree topology in a switched network). For each VLAN, the switch with the highest switch priority (the lowest numerical priority value) is elected as the root switch. If all switches are configured with the default priority (32768), the switch with the lowest MAC address in the VLAN becomes the root switch. The switch priority value occupies the most significant bits of the bridge ID, as shown in [Table 14-1 on page 14-4](#).

- A root port is selected for each switch (except the root switch). This port provides the best path (lowest cost) when the switch forwards packets to the root switch.
- The shortest distance to the root switch is calculated for each switch based on the path cost.
- A designated switch for each LAN segment is selected. The designated switch incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.

All paths that are not needed to reach the root switch from anywhere in the switched network are placed in the spanning-tree blocking mode.

Bridge ID, Switch Priority, and Extended System ID

The IEEE 802.1D standard requires that each switch has a unique bridge identifier (bridge ID), which determines the selection of the root switch. Because each VLAN is considered as a different *logical bridge* with PVST+ and rapid PVST+, the same switch must have as many different bridge IDs as VLANs configured on it. Each VLAN on the switch has a unique 8-byte bridge ID. The two most-significant bytes are used for the switch priority, and the remaining six bytes are derived from the switch MAC address.

The Catalyst 2970 switch supports the 802.1T spanning-tree extensions, and some of the bits previously used for the switch priority are now used as the VLAN identifier. The result is that fewer MAC addresses are reserved for the switch, and a larger range of VLAN IDs can be supported, all while maintaining the uniqueness of the bridge ID. As shown in [Table 14-1](#), the two bytes previously used for the switch priority are reallocated into a 4-bit priority value and a 12-bit extended system ID value equal to the VLAN ID.

Table 14-1 Switch Priority Value and Extended System ID

Switch Priority Value				Extended System ID (Set Equal to the VLAN ID)											
Bit 16	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1
32768	16384	8192	4096	2048	1024	512	256	128	64	32	16	8	4	2	1

Spanning tree uses the extended system ID, the switch priority, and the allocated spanning-tree MAC address to make the bridge ID unique for each VLAN.

Support for the extended system ID affects how you manually configure the root switch, the secondary root switch, and the switch priority of a VLAN. For example, when you change the switch priority value, you change the probability that the switch will be elected as the root switch. Configuring a higher value decreases the probability; a lower value increases the probability. For more information, see the [“Configuring the Root Switch”](#) section on page 14-14, the [“Configuring a Secondary Root Switch”](#) section on page 14-16, and the [“Configuring the Switch Priority of a VLAN”](#) section on page 14-19.

Spanning-Tree Interface States

Propagation delays can occur when protocol information passes through a switched LAN. As a result, topology changes can take place at different times and at different places in a switched network. When an interface transitions directly from nonparticipation in the spanning-tree topology to the forwarding state, it can create temporary data loops. Interfaces must wait for new topology information to propagate through the switched LAN before starting to forward frames. They must allow the frame lifetime to expire for forwarded frames that have used the old topology.

Each Layer 2 interface on a switch using spanning tree exists in one of these states:

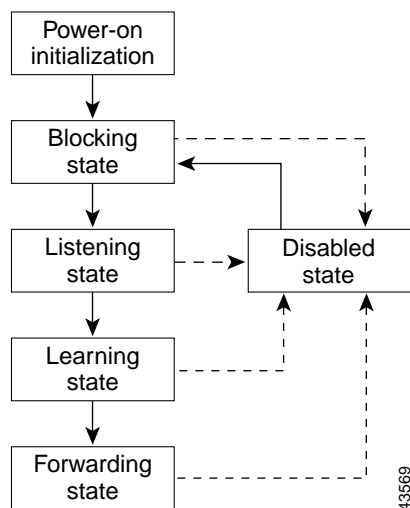
- **Blocking**—The interface does not participate in frame forwarding.
- **Listening**—The first transitional state after the blocking state when the spanning tree determines that the interface should participate in frame forwarding.
- **Learning**—The interface prepares to participate in frame forwarding.
- **Forwarding**—The interface forwards frames.
- **Disabled**—The interface is not participating in spanning tree because of a shutdown port, no link on the port, or no spanning-tree instance running on the port.

An interface moves through these states:

- From initialization to blocking
- From blocking to listening or to disabled
- From listening to learning or to disabled
- From learning to forwarding or to disabled
- From forwarding to disabled

Figure 14-1 illustrates how an interface moves through the states.

Figure 14-1 Spanning-Tree Interface States



When you power up the switch, spanning tree is enabled by default, and every interface in the switch, VLAN, or network goes through the blocking state and the transitory states of listening and learning. Spanning tree stabilizes each interface at the forwarding or blocking state.

When the spanning-tree algorithm places a Layer 2 interface in the forwarding state, this process occurs:

1. The interface is in the listening state while spanning tree waits for protocol information to transition the interface to the blocking state.
2. While spanning tree waits the forward-delay timer to expire, it moves the interface to the learning state and resets the forward-delay timer.
3. In the learning state, the interface continues to block frame forwarding as the switch learns end-station location information for the forwarding database.
4. When the forward-delay timer expires, spanning tree moves the interface to the forwarding state, where both learning and frame forwarding are enabled.

Blocking State

A Layer 2 interface in the blocking state does not participate in frame forwarding. After initialization, a BPDU is sent to each switch interface. A switch initially functions as the root until it exchanges BPDUs with other switches. This exchange establishes which switch in the network is the root or root switch. If

there is only one switch in the network, no exchange occurs, the forward-delay timer expires, and the interfaces move to the listening state. An interface always enters the blocking state after switch initialization.

An interface in the blocking state performs these functions:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Listening State

The listening state is the first state a Layer 2 interface enters after the blocking state. The interface enters this state when the spanning tree determines that the interface should participate in frame forwarding.

An interface in the listening state performs these functions:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Receives BPDUs

Learning State

A Layer 2 interface in the learning state prepares to participate in frame forwarding. The interface enters the learning state from the listening state.

An interface in the learning state performs these functions:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Learns addresses
- Receives BPDUs

Forwarding State

A Layer 2 interface in the forwarding state forwards frames. The interface enters the forwarding state from the learning state.

An interface in the forwarding state performs these functions:

- Receives and forwards frames received on the port
- Forwards frames switched from another port
- Learns addresses
- Receives BPDUs

Disabled State

A Layer 2 interface in the disabled state does not participate in frame forwarding or in the spanning tree. An interface in the disabled state is nonoperational.

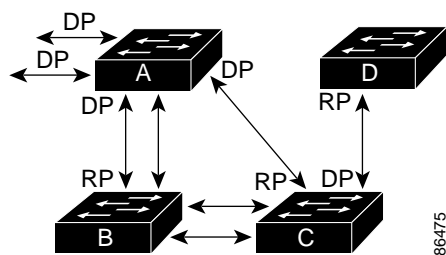
A disabled interface performs these functions:

- Discards frames received on the port
- Discards frames switched from another interface for forwarding
- Does not learn addresses
- Does not receive BPDUs

How a Switch or Port Becomes the Root Switch or Root Port

If all switches in a network are enabled with default spanning-tree settings, the switch with the lowest MAC address becomes the root switch. In [Figure 14-2](#), Switch A is elected as the root switch because the switch priority of all the switches is set to the default (32768) and Switch A has the lowest MAC address. However, because of traffic patterns, number of forwarding interfaces, or link types, Switch A might not be the ideal root switch. By increasing the priority (lowering the numerical value) of the ideal switch so that it becomes the root switch, you force a spanning-tree recalculation to form a new topology with the ideal switch as the root.

Figure 14-2 Spanning-Tree Topology



RP = Root Port
DP = Designated Port

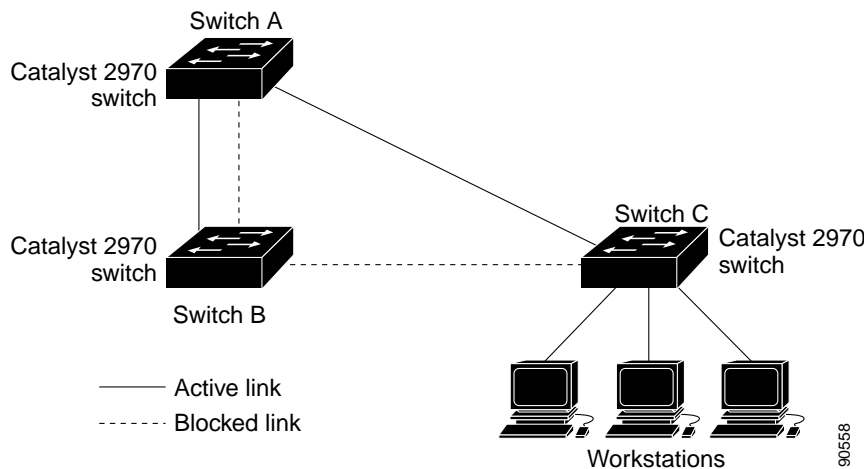
When the spanning-tree topology is calculated based on default parameters, the path between source and destination end stations in a switched network might not be ideal. For instance, connecting higher-speed links to an interface that has a higher number than the root port can cause a root-port change. The goal is to make the fastest link the root port.

For example, assume that one port on Switch B is a Gigabit Ethernet link and that another port on Switch B (a 10/100 link) is the root port. Network traffic might be more efficient over the Gigabit Ethernet link. By changing the spanning-tree port priority on the Gigabit Ethernet interface to a higher priority (lower numerical value) than the root port, the Gigabit Ethernet interface becomes the new root port.

Spanning Tree and Redundant Connectivity

You can create a redundant backbone with spanning tree by connecting two switch interfaces to another device or to two different devices, as shown in [Figure 14-3](#). Spanning tree automatically disables one interface but enables it if the other one fails. If one link is high-speed and the other is low-speed, the low-speed link is always disabled. If the speeds are the same, the port priority and port ID are added together, and spanning tree disables the link with the lowest value.

Figure 14-3 Spanning Tree and Redundant Connectivity



You can also create redundant links between switches by using EtherChannel groups. For more information, see the [Chapter 28, “Configuring EtherChannels.”](#)

Spanning-Tree Address Management

IEEE 802.1D specifies 17 multicast addresses, ranging from 0x0180C2000000 to 0x0180C2000010, to be used by different bridge protocols. These addresses are static addresses that cannot be removed.

Regardless of the spanning-tree state, each switch receives but does not forward packets destined for addresses between 0x0180C2000000 and 0x0180C200000F.

If spanning tree is enabled, the CPU on the switch receives packets destined for 0x0180C2000000 and 0x0180C2000010. If spanning tree is disabled, the switch forwards those packets as unknown multicast addresses.

Accelerated Aging to Retain Connectivity

The default for aging dynamic addresses is 5 minutes, the default setting of the **mac address-table aging-time** global configuration command. However, a spanning-tree reconfiguration can cause many station locations to change. Because these stations could be unreachable for 5 minutes or more during a reconfiguration, the address-aging time is accelerated so that station addresses can be dropped from the address table and then relearned. The accelerated aging is the same as the forward-delay parameter value (**spanning-tree vlan *vlan-id* forward-time *seconds*** global configuration command) when the spanning tree reconfigures.

Because each VLAN is a separate spanning-tree instance, the switch accelerates aging on a per-VLAN basis. A spanning-tree reconfiguration on one VLAN can cause the dynamic addresses learned on that VLAN to be subject to accelerated aging. Dynamic addresses on other VLANs can be unaffected and remain subject to the aging interval entered for the switch.

Spanning-Tree Modes and Protocols

The switch supports these spanning-tree modes and protocols:

- **PVST+—**This spanning-tree mode is based on the IEEE 802.1D standard and Cisco proprietary extensions. It is the default spanning-tree mode used on all Gigabit Ethernet port-based VLANs. The PVST+ runs on each VLAN on the switch up to the maximum supported, ensuring that each has a loop-free path through the network.

The PVST+ provides Layer 2 load balancing for the VLAN on which it runs. You can create different logical topologies by using the VLANs on your network to ensure that all of your links are used but that no one link is oversubscribed. Each instance of PVST+ on a VLAN has a single root switch. This root switch propagates the spanning-tree information associated with that VLAN to all other switches in the network. Because each switch has the same information about the network, this process ensures that the network topology is maintained.

- **Rapid PVST+—**This spanning-tree mode is the same as PVST+ except that it uses a rapid convergence based on the IEEE 802.1W standard. To provide rapid convergence, the rapid PVST+ immediately deletes dynamically learned MAC address entries on a per-port basis upon receiving a topology change. By contrast, PVST+ uses a short aging time for dynamically learned MAC address entries.

The rapid PVST+ uses the same configuration as PVST+ (except where noted), and the switch needs only minimal extra configuration. The benefit of rapid PVST+ is that you can migrate a large PVST+ install base to rapid PVST+ without having to learn the complexities of the MSTP configuration and without having to reprovision your network. In rapid-PVST+ mode, each VLAN runs its own spanning-tree instance up to the maximum supported.

- **MSTP—**This spanning-tree mode is based on the IEEE 802.1S standard. You can map multiple VLANs to the same spanning-tree instance, which reduces the number of spanning-tree instances required to support a large number of VLANs. The MSTP runs on top of the RSTP (based on IEEE 802.1W), which provides for rapid convergence of the spanning tree by eliminating the forward delay and by quickly transitioning root ports and designated ports to the forwarding state. You cannot run MSTP without RSTP.

The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network. For more information, see [Chapter 15, “Configuring MSTP.”](#)

For information about the number of supported spanning-tree instances, see the next section.

Supported Spanning-Tree Instances

In PVST+ or rapid-PVST+ mode, the switch supports up to 128 spanning-tree instances.

In MSTP mode, the switch supports up to 16 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.

For information about how spanning tree interoperates with the VLAN Trunking Protocol (VTP), see the [“Spanning-Tree Configuration Guidelines”](#) section on page 14-12.

Spanning-Tree Interoperability and Backward Compatibility

[Table 14-2](#) lists the interoperability and compatibility among the supported spanning-tree modes in a network.

Table 14-2 *PVST+, MSTP, and Rapid-PVST+ Interoperability*

	PVST+	MSTP	Rapid PVST+
PVST+	Yes	Yes (with restrictions)	Yes (reverts to PVST+)
MSTP	Yes (with restrictions)	Yes	Yes (reverts to PVST+)
Rapid PVST+	Yes (reverts to PVST+)	Yes (reverts to PVST+)	Yes

In a mixed MSTP and PVST+ network, the common spanning-tree (CST) root must be inside the MST backbone, and a PVST+ switch cannot connect to multiple MST regions.

When a network contains switches running rapid PVST+ and switches running PVST+, we recommend that the rapid-PVST+ switches and PVST+ switches be configured for different spanning-tree instances. In the rapid-PVST+ spanning-tree instances, the root switch must be a rapid-PVST+ switch. In the PVST+ instances, the root switch must be a PVST+ switch. The PVST+ switches should be at the edge of the network.

STP and IEEE 802.1Q Trunks

The IEEE 802.1Q standard for VLAN trunks imposes some limitations on the spanning-tree strategy for a network. The standard requires only one spanning-tree instance for *all* VLANs allowed on the trunks. However, in a network of Cisco switches connected through 802.1Q trunks, the switches maintain one spanning-tree instance for *each* VLAN allowed on the trunks.

When you connect a Cisco switch to a non-Cisco device through an 802.1Q trunk, the Cisco switch uses PVST+ to provide spanning-tree interoperability. If rapid PVST+ is enabled, the switch uses it instead of PVST+. The switch combines the spanning-tree instance of the 802.1Q VLAN of the trunk with the spanning-tree instance of the non-Cisco 802.1Q switch.

However, all PVST+ or rapid-PVST+ information is maintained by Cisco switches separated by a cloud of non-Cisco 802.1Q switches. The non-Cisco 802.1Q cloud separating the Cisco switches is treated as a single trunk link between the switches.

PVST+ is automatically enabled on 802.1Q trunks, and no user configuration is required. The external spanning-tree behavior on access ports and Inter-Switch Link (ISL) trunk ports is not affected by PVST+.

For more information on 802.1Q trunks, see [Chapter 11, “Configuring VLANs.”](#)

Configuring Spanning-Tree Features

These sections describe how to configure spanning-tree features:

- [Default Spanning-Tree Configuration, page 14-11](#)
- [Spanning-Tree Configuration Guidelines, page 14-12](#)
- [Changing the Spanning-Tree Mode, page 14-13](#) (required)
- [Disabling Spanning Tree, page 14-14](#) (optional)
- [Configuring the Root Switch, page 14-14](#) (optional)
- [Configuring a Secondary Root Switch, page 14-16](#) (optional)
- [Configuring Port Priority, page 14-17](#) (optional)
- [Configuring Path Cost, page 14-18](#) (optional)
- [Configuring the Switch Priority of a VLAN, page 14-19](#) (optional)
- [Configuring Spanning-Tree Timers, page 14-20](#) (optional)

Default Spanning-Tree Configuration

[Table 14-3](#) shows the default spanning-tree configuration.

Table 14-3 Default Spanning-Tree Configuration

Feature	Default Setting
Enable state	Enabled on VLAN 1. For more information, see the “Supported Spanning-Tree Instances” section on page 14-9 .
Spanning-tree mode	PVST+. (Rapid PVST+ and MSTP are disabled.)
Switch priority	32768.
Spanning-tree port priority (configurable on a per-interface basis)	128.
Spanning-tree port cost (configurable on a per-interface basis)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Spanning-tree VLAN port priority (configurable on a per-VLAN basis)	128.
Spanning-tree VLAN port cost (configurable on a per-VLAN basis)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Spanning-tree timers	Hello time: 2 seconds. Forward-delay time: 15 seconds. Maximum-aging time: 20 seconds.

Spanning-Tree Configuration Guidelines

If more VLANs are defined in the VTP than there are spanning-tree instances, you can enable PVST+ or rapid PVST+ on only 128 VLANs on the switch. The remaining VLANs operate with spanning tree disabled. However, you can map multiple VLANs to the same spanning-tree instances by using MSTP. For more information, see [Chapter 15, “Configuring MSTP.”](#)

If 128 instances of spanning tree are already in use, you can disable spanning tree on one of the VLANs and then enable it on the VLAN where you want it to run. Use the **no spanning-tree vlan *vlan-id*** global configuration command to disable spanning tree on a specific VLAN, and use the **spanning-tree vlan *vlan-id*** global configuration command to enable spanning tree on the desired VLAN.



Caution

Switches that are not running spanning tree still forward BPDUs that they receive so that the other switches on the VLAN that have a running spanning-tree instance can break loops. Therefore, spanning tree must be running on enough switches to break all the loops in the network; for example, at least one switch on each loop in the VLAN must be running spanning tree. It is not absolutely necessary to run spanning tree on all switches in the VLAN. However, if you are running spanning tree only on a minimal set of switches, an incautious change to the network that introduces another loop into the VLAN can result in a broadcast storm.



Note

If you have already used all available spanning-tree instances on your switch, adding another VLAN anywhere in the VTP domain creates a VLAN that is not running spanning tree on that switch. If you have the default allowed list on the trunk ports of that switch, the new VLAN is carried on all trunk ports. Depending on the topology of the network, this could create a loop in the new VLAN that will not be broken, particularly if there are several adjacent switches that have all run out of spanning-tree instances. You can prevent this possibility by setting up allowed lists on the trunk ports of switches that have used up their allocation of spanning-tree instances. Setting up allowed lists is not necessary in many cases and can make it more labor-intensive to add another VLAN to the network.

Spanning-tree commands determine the configuration of VLAN spanning-tree instances. You create a spanning-tree instance when you assign an interface to a VLAN. The spanning-tree instance is removed when the last interface is moved to another VLAN. You can configure switch and port parameters before a spanning-tree instance is created; these parameters are applied when the spanning-tree instance is created.

The switch supports PVST+, rapid PVST+, and MSTP, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For information about the different spanning-tree modes and how they interoperate, see the [“Spanning-Tree Interoperability and Backward Compatibility” section on page 14-10](#).

The UplinkFast and BackboneFast features are not supported with the rapid PVST+.

Changing the Spanning-Tree Mode

The switch supports three spanning-tree modes: PVST+, rapid PVST+, or MSTP. By default, the switch runs the PVST+ protocol.

Beginning in privileged EXEC mode, follow these steps to change the spanning-tree mode. If you want to enable a mode that is different from the default mode, this procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mode {pvst mst rapid-pvst}	Configure a spanning-tree mode. <ul style="list-style-type: none"> • Select pvst to enable PVST+ (the default setting). • Select mst to enable MSTP (and RSTP). For more configuration steps, see Chapter 15, “Configuring MSTP.” • Select rapid-pvst to enable rapid PVST+.
Step 3	interface <i>interface-id</i>	(Recommended for rapid-PVST+ mode only) Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical ports, VLANs, and port channels. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 12.
Step 4	spanning-tree link-type point-to-point	(Recommended for rapid-PVST+ mode only) Specify that the link type for this port is point-to-point. If you connect this port (local port) to a remote port through a point-to-point link and the local port becomes a designated port, the switch negotiates with the remote port and rapidly transitions the local port to the forwarding state.
Step 5	end	Return to privileged EXEC mode.
Step 6	clear spanning-tree detected-protocols	(Recommended for rapid-PVST+ mode only) If any port on the switch is connected to a port on a legacy 802.1D switch, restart the protocol migration process on the entire switch. This step is optional if the designated switch determines that this switch is running rapid PVST+.
Step 7	show spanning-tree summary and show spanning-tree interface <i>interface-id</i>	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree mode** global configuration command. To return the port to its default setting, use the **no spanning-tree link-type** interface configuration command.

Disabling Spanning Tree

Spanning tree is enabled by default on VLAN 1 and on all newly created VLANs up to the spanning-tree limit specified in the [“Supported Spanning-Tree Instances” section on page 14-9](#). Disable spanning tree only if you are sure there are no loops in the network topology.



Caution

When spanning tree is disabled and loops are present in the topology, excessive traffic and indefinite packet duplication can drastically reduce network performance.

Beginning in privileged EXEC mode, follow these steps to disable spanning-tree on a per-VLAN basis. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no spanning-tree vlan <i>vlan-id</i>	For <i>vlan-id</i> , the range is 1 to 4094.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable spanning-tree, use the **spanning-tree vlan** *vlan-id* global configuration command.

Configuring the Root Switch

The switch maintains a separate spanning-tree instance for each active VLAN configured on it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For each VLAN, the switch with the lowest bridge ID becomes the root switch for that VLAN.

To configure a switch to become the root for the specified VLAN, use the **spanning-tree vlan** *vlan-id* **root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value. When you enter this command, the software checks the switch priority of the root switches for each VLAN. Because of the extended system ID support, the switch sets its own priority for the specified VLAN to 24576 if this value will cause this switch to become the root for the specified VLAN.

If any root switch for the specified VLAN has a switch priority lower than 24576, the switch sets its own priority for the specified VLAN to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in [Table 14-1 on page 14-4](#).)



Note

The **spanning-tree vlan** *vlan-id* **root** global configuration command fails if the value necessary to be the root switch is less than 1.



Note

If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

**Note**

The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

**Note**

After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time through the **spanning-tree vlan *vlan-id* hello-time**, **spanning-tree vlan *vlan-id* forward-time**, and the **spanning-tree vlan *vlan-id* max-age** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to configure a switch to become the root for the specified VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> root primary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch to become the root for the specified VLAN. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10; the default is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree detail	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring a Secondary Root Switch

When you configure a Catalyst 2970 switch as the secondary root, the switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified VLAN if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values that you used when you configured the primary root switch with the **spanning-tree vlan *vlan-id* root primary** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a switch to become the secondary root for the specified VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch to become the secondary root for the specified VLAN. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10; the default is 2. Use the same network diameter and hello-time values that you used when configuring the primary root switch. See the “Configuring the Root Switch” section on page 14-14 .
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree detail	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* root** global configuration command.

Configuring Port Priority

If a loop occurs, spanning tree uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the port priority of an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>).
Step 3	spanning-tree port-priority <i>priority</i>	Configure the port priority for an interface. For <i>priority</i> , the range is 0 to 240, in increments of 16; the default is 128. Valid values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected. The lower the number, the higher the priority.
Step 4	spanning-tree vlan <i>vlan-id</i> port-priority <i>priority</i>	Configure the port priority for a VLAN. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>priority</i>, the range is 0 to 240, in increments of 16; the default is 128. Valid values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected. The lower the number, the higher the priority.
Step 5	end	Return to privileged EXEC mode.
Step 6	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only if the port is in a link-up operative state. Otherwise, you can use the **show running-config interface** privileged EXEC command to confirm the configuration.

To return to the default setting, use the **no spanning-tree [vlan *vlan-id*] port-priority** interface configuration command. For information on how to configure load sharing on trunk ports by using spanning-tree port priorities, see the [“Configuring Trunk Ports for Load Sharing”](#) section on page 11-22.

Configuring Path Cost

The spanning-tree path cost default value is derived from the media speed of an interface. If a loop occurs, spanning tree uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, spanning tree puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the cost of an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>).
Step 3	spanning-tree cost <i>cost</i>	Configure the cost for an interface. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. For <i>cost</i> , the range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 4	spanning-tree vlan <i>vlan-id</i> cost <i>cost</i>	Configure the cost for a VLAN. If a loop occurs, spanning tree uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 5	end	Return to privileged EXEC mode.
Step 6	show spanning-tree interface <i>interface-id</i> or show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

**Note**

The **show spanning-tree interface** *interface-id* privileged EXEC command displays information only for ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return to the default setting, use the **no spanning-tree [vlan *vlan-id*] cost** interface configuration command. For information on how to configure load sharing on trunk ports by using spanning-tree path costs, see the [“Configuring Trunk Ports for Load Sharing”](#) section on page 11-22.

Configuring the Switch Priority of a VLAN

You can configure the switch priority and make it more likely that the switch will be chosen as the root switch.

**Note**

Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the switch priority.

Beginning in privileged EXEC mode, follow these steps to configure the switch priority of a VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> priority <i>priority</i>	Configure the switch priority of a VLAN. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>priority</i>, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Valid priority values are 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* priority** global configuration command.

Configuring Spanning-Tree Timers

Table 14-4 describes the timers that affect the entire spanning-tree performance.

Table 14-4 *Spanning-Tree Timers*

Variable	Description
Hello timer	Determines how often the switch broadcasts hello messages to other switches.
Forward-delay timer	Determines how long each of the listening and learning states last before the interface begins forwarding.
Maximum-age timer	Determines the amount of time the switch stores protocol information received on an interface.

The sections that follow provide the configuration steps.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root switch by changing the hello time.



Note

Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree vlan *vlan-id* root primary** and the **spanning-tree vlan *vlan-id* root secondary** global configuration commands to modify the hello time.

Beginning in privileged EXEC mode, follow these steps to configure the hello time of a VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> hello-time <i>seconds</i>	Configure the hello time of a VLAN. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>seconds</i>, the range is 1 to 10; the default is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* hello-time** global configuration command.

Configuring the Forwarding-Delay Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the forwarding-delay time for a VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> forward-time <i>seconds</i>	Configure the forward time of a VLAN. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>seconds</i>, the range is 4 to 30; the default is 15.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* forward-time** global configuration command.

Configuring the Maximum-Aging Time for a VLAN

Beginning in privileged EXEC mode, follow these steps to configure the maximum-aging time for a VLAN. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree vlan <i>vlan-id</i> max-age <i>seconds</i>	Configure the maximum-aging time of a VLAN. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. <ul style="list-style-type: none"> For <i>vlan-id</i>, you can specify a single VLAN identified by VLAN ID number, a range of VLANs separated by a hyphen, or a series of VLANs separated by a comma. The range is 1 to 4094. For <i>seconds</i>, the range is 6 to 40; the default is 20.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree vlan <i>vlan-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no spanning-tree vlan *vlan-id* max-age** global configuration command.

Displaying the Spanning-Tree Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in [Table 14-5](#):

Table 14-5 *Commands for Displaying Spanning-Tree Status*

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface <i>interface-id</i>	Displays spanning-tree information for the specified interface.
show spanning-tree summary [totals]	Displays a summary of port states or displays the total lines of the STP state section.

You can clear spanning-tree counters by using the **clear spanning-tree [interface *interface-id*]** privileged EXEC command.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.



Configuring MSTP

This chapter describes how to configure the Cisco implementation of the IEEE 802.1S Multiple STP (MSTP) on the Catalyst 2970 switch.

The MSTP enables multiple VLANs to be mapped to the same spanning-tree instance, thereby reducing the number of spanning-tree instances needed to support a large number of VLANs. The MSTP provides for multiple forwarding paths for data traffic and enables load balancing. It improves the fault tolerance of the network because a failure in one instance (forwarding path) does not affect other instances (forwarding paths). The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network; this deployment provides the highly available network required in a service-provider environment.

When the switch is in the multiple spanning-tree (MST) mode, the Rapid Spanning Tree Protocol (RSTP), which is based on IEEE 802.1W, is automatically enabled. The RSTP provides rapid convergence of the spanning tree through explicit handshaking that eliminates the IEEE 802.1D forwarding delay and quickly transitions root ports and designated ports to the forwarding state.

Both MSTP and RSTP improve the spanning-tree operation and maintain backward compatibility with equipment that is based on the (original) 802.1D spanning tree, with existing Cisco-proprietary Multiple Instance STP (MISTP), and with existing Cisco per-VLAN spanning-tree plus (PVST+) and rapid per-VLAN spanning-tree plus (rapid PVST+). For information about PVST+ and rapid PVST+, see [Chapter 14, “Configuring STP.”](#) For information about other spanning-tree features such as Port Fast, UplinkFast, root guard, and so forth, see [Chapter 16, “Configuring Optional Spanning-Tree Features.”](#)



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding MSTP, page 15-2](#)
- [Understanding RSTP, page 15-6](#)
- [Configuring MSTP Features, page 15-11](#)
- [Displaying the MST Configuration and Status, page 15-23](#)

Understanding MSTP

MSTP, which uses RSTP for rapid convergence, enables VLANs to be grouped into a spanning-tree instance, with each instance having a spanning-tree topology independent of other spanning-tree instances. This architecture provides multiple forwarding paths for data traffic, enables load balancing, and reduces the number of spanning-tree instances required to support a large number of VLANs.

These sections describe how the MSTP works:

- [Multiple Spanning-Tree Regions, page 15-2](#)
- [IST, CIST, and CST, page 15-3](#)
- [Hop Count, page 15-5](#)
- [Boundary Ports, page 15-5](#)
- [“Interoperability with 802.1D STP” section on page 15-5](#)

For configuration information, see the [“Configuring MSTP Features” section on page 15-11](#).

Multiple Spanning-Tree Regions

For switches to participate in multiple spanning-tree (MST) instances, you must consistently configure the switches with the same MST configuration information. A collection of interconnected switches that have the same MST configuration comprises an MST region as shown in [Figure 15-1 on page 15-4](#).

The MST configuration determines to which MST region each switch belongs. The configuration includes the name of the region, the revision number, and the MST VLAN-to-instance assignment map. You configure the switch for a region by using the **spanning-tree mst configuration** global configuration command, after which the switch enters the MST configuration mode. From this mode, you can map VLANs to an MST instance by using the **instance** MST configuration command, specify the region name by using the **name** MST configuration command, and set the revision number by using the **revision** MST configuration command.

A region can have one member or multiple members with the same MST configuration; each member must be capable of processing RSTP bridge protocol data units (BPDUs). There is no limit to the number of MST regions in a network, but each region can support up to 16 spanning-tree instances. You can assign a VLAN to only one spanning-tree instance at a time.

IST, CIST, and CST

Unlike PVST+ and rapid PVST+ in which all the spanning-tree instances are independent, the MSTP establishes and maintains two types of spanning trees:

- An internal spanning tree (IST), which is the spanning tree that runs in an MST region.

Within each MST region, the MSTP maintains multiple spanning-tree instances. Instance 0 is a special instance for a region, known as the internal spanning tree (IST). All other MST instances are numbered from 1 to 15.

The IST is the only spanning-tree instance that sends and receives BPDUs; all of the other spanning-tree instance information is contained in M-records, which are encapsulated within MSTP BPDUs. Because the MSTP BPDU carries information for all instances, the number of BPDUs that need to be processed by a switch to support multiple spanning-tree instances is significantly reduced.

All MST instances within the same region share the same protocol timers, but each MST instance has its own topology parameters, such as root switch ID, root path cost, and so forth. By default, all VLANs are assigned to the IST.

An MST instance is local to the region; for example, MST instance 1 in region A is independent of MST instance 1 in region B, even if regions A and B are interconnected.

- A common and internal spanning tree (CIST), which is a collection of the ISTs in each MST region, and the common spanning tree (CST) that interconnects the MST regions and single spanning trees.

The spanning tree computed in a region appears as a subtree in the CST that encompasses the entire switched domain. The CIST is formed as a result of the spanning-tree algorithm running between switches that support the 802.1W, 802.1S, and 802.1D protocols. The CIST inside an MST region is the same as the CST outside a region.

For more information, see the [“Operations Within an MST Region” section on page 15-3](#) and the [“Operations Between MST Regions” section on page 15-4](#).

Operations Within an MST Region

The IST connects all the MSTP switches in a region. When the IST converges, the root of the IST becomes the IST master (shown in [Figure 15-1 on page 15-4](#)), which is the switch within the region with the lowest bridge ID and path cost to the CST root. The IST master also is the CST root if there is only one region within the network. If the CST root is outside the region, one of the MSTP switches at the boundary of the region is selected as the IST master.

When an MSTP switch initializes, it sends BPDUs claiming itself as the root of the CST and the IST master, with both of the path costs to the CST root and to the IST master set to zero. The switch also initializes all of its MST instances and claims to be the root for all of them. If the switch receives superior MST root information (lower bridge ID, lower path cost, and so forth) than currently stored for the port, it relinquishes its claim as the IST master.

During initialization, a region might have many subregions, each with its own IST master. As switches receive superior IST information, they leave their old subregions and join the new subregion that might contain the true IST master. Thus all subregions shrink, except for the one that contains the true IST master.

For correct operation, all switches in the MST region must agree on the same IST master. Therefore, any two switches in the region synchronize their port roles for an MST instance only if they converge to a common IST master.

Operations Between MST Regions

If there are multiple regions or legacy 802.1D switches within the network, MSTP establishes and maintains the CST, which includes all MST regions and all legacy STP switches in the network. The MST instances combine with the IST at the boundary of the region to become the CST.

The IST connects all the MSTP switches in the region and appears as a subtree in the CST that encompasses the entire switched domain, with the root of the subtree being the IST master. The MST region appears as a virtual switch to adjacent STP switches and MST regions.

Figure 15-1 shows a network with three MST regions and a legacy 802.1D switch (D). The IST master for region 1 (A) is also the CST root. The IST master for region 2 (B) and the IST master for region 3 (C) are the roots for their respective subtrees within the CST. The RSTP runs in all regions.

Figure 15-1 MST Regions, IST Masters, and the CST Root

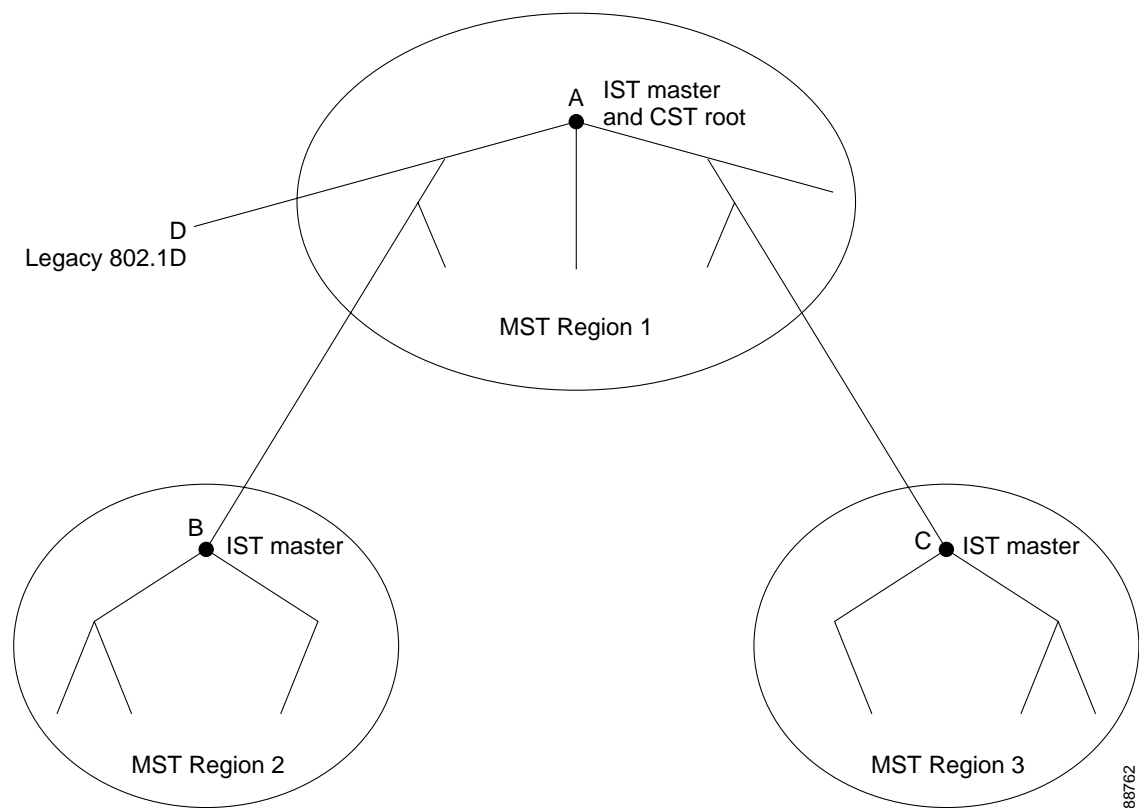


Figure 15-1 does not show additional MST instances for each region. Note that the topology of MST instances can be different from that of the IST for the same region.

Only the CST instance sends and receives BPDUs, and MST instances add their spanning-tree information into the BPDUs to interact with neighboring switches and compute the final spanning-tree topology. Because of this, the spanning-tree parameters related to BPDU transmission (for example, hello time, forward time, max-age, and max-hops) are configured only on the CST instance but affect all MST instances. Parameters related to the spanning-tree topology (for example, switch priority, port VLAN cost, port VLAN priority) can be configured on both the CST instance and the MST instance.

MSTP switches use version 3 RSTP BPDUs or 802.1D STP BPDUs to communicate with legacy 802.1D switches. MSTP switches use MSTP BPDUs to communicate with MSTP switches.

Hop Count

The IST and MST instances do not use the message-age and maximum-age information in the configuration BPDU to compute the spanning-tree topology. Instead, they use the path cost to the root and a hop-count mechanism similar to the IP time-to-live (TTL) mechanism.

By using the **spanning-tree mst max-hops** global configuration command, you can configure the maximum hops inside the region and apply it to the IST and all MST instances in that region. The hop count achieves the same result as the message-age information (determines when to trigger a reconfiguration). The root switch of the instance always sends a BPDU (or M-record) with a cost of 0 and the hop count set to the maximum value. When a switch receives this BPDU, it decrements the received remaining hop count by one and propagates this value as the remaining hop count in the BPDUs it generates. When the count reaches zero, the switch discards the BPDU and ages the information held for the port.

The message-age and maximum-age information in the RSTP portion of the BPDU remain the same throughout the region, and the same values are propagated by the region's designated ports at the boundary.

Boundary Ports

A boundary port is a port that connects an MST region to a single spanning-tree region running RSTP, to a single spanning-tree region running PVST+ or rapid PVST+, or to another MST region with a different MST configuration. A boundary port also connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

At the boundary, the roles of the MST ports do not matter, and their state is forced to be the same as the IST port state (MST ports at the boundary are in the forwarding state only when the IST port is forwarding). An IST port at the boundary can have any port role except a backup port role.

On a shared boundary link, the MST ports wait in the blocking state for the forward-delay time to expire before transitioning to the learning state. The MST ports wait another forward-delay time before transitioning to the forwarding state.

If the boundary port is on a point-to-point link and it is the IST root port, the MST ports transition to the forwarding state as soon as the IST port transitions to the forwarding state.

If the IST port is a designated port on a point-to-point link and if the IST port transitions to the forwarding state because of an agreement received from its peer port, the MST ports also immediately transition to the forwarding state.

If a boundary port transitions to the forwarding state in an IST instance, it is forwarding in all MST instances, and a topology change is triggered. If a boundary port with the IST root or designated port role receives a topology change notice external to the MST cloud, the MSTP switch triggers a topology change in the IST instance and in all the MST instances active on that port.

Interoperability with 802.1D STP

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy 802.1D switches. If this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MSTP switch also can detect that a port is at the boundary of a region when it receives a legacy BPDU, an MSTP BPDU (version 3) associated with a different region, or an RSTP BPDU (version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives 802.1D BPDUs because it cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. Also, a switch might continue to assign a boundary role to a port when the switch to which this switch is connected has joined the region. To restart the protocol migration process (force the renegotiation with neighboring switches), use the **clear spanning-tree detected-protocols** privileged EXEC command.

If all the legacy switches on the link are RSTP switches, they can process MSTP BPDUs as if they are RSTP BPDUs. Therefore, MSTP switches send either a version 0 configuration and TCN BPDUs or version 3 MSTP BPDUs on a boundary port. A boundary port connects to a LAN, the designated switch of which is either a single spanning-tree switch or a switch with a different MST configuration.

Understanding RSTP

The RSTP takes advantage of point-to-point wiring and provides rapid convergence of the spanning tree. Reconfiguration of the spanning tree can occur in less than 1 second (in contrast to 50 seconds with the default settings in the 802.1D spanning tree), which is critical for networks carrying delay-sensitive traffic such as voice and video.

This section describes how the RSTP works:

- [Port Roles and the Active Topology, page 15-6](#)
- [Rapid Convergence, page 15-7](#)
- [Synchronization of Port Roles, page 15-8](#)
- [Bridge Protocol Data Unit Format and Processing, page 15-9](#)

For configuration information, see the “[Configuring MSTP Features](#)” section on [page 15-11](#).

Port Roles and the Active Topology

The RSTP provides rapid convergence of the spanning tree by assigning port roles and by determining the active topology. The RSTP builds upon the IEEE 802.1D STP to select the switch with the highest switch priority (lowest numerical priority value) as the root switch as described in the “[Spanning-Tree Topology and BPDUs](#)” section on [page 14-3](#). Then the RSTP assigns one of these port roles to individual ports:

- Root port—Provides the best path (lowest cost) when the switch forwards packets to the root switch.
- Designated port—Connects to the designated switch, which incurs the lowest path cost when forwarding packets from that LAN to the root switch. The port through which the designated switch is attached to the LAN is called the designated port.
- Alternate port—Offers an alternate path toward the root switch to that provided by the current root port.
- Backup port—Acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. A backup port can exist only when two ports are connected together in a loopback by a point-to-point link or when a switch has two or more connections to a shared LAN segment.
- Disabled port—Has no role within the operation of the spanning tree.

A port with the root or a designated port role is included in the active topology. A port with the alternate or backup port role is excluded from the active topology.

In a stable topology with consistent port roles throughout the network, the RSTP ensures that every root port and designated port immediately transition to the forwarding state while all alternate and backup ports are always in the discarding state (equivalent to blocking in 802.1D). The port state controls the operation of the forwarding and learning processes. Table 15-1 provides a comparison of 802.1D and RSTP port states.

Table 15-1 Port State Comparison

Operational Status	STP Port State (802.1D)	RSTP Port State	Is Port Included in the Active Topology?
Enabled	Blocking	Discarding	No
Enabled	Listening	Discarding	No
Enabled	Learning	Learning	Yes
Enabled	Forwarding	Forwarding	Yes
Disabled	Disabled	Discarding	No

To be consistent with Cisco STP implementations, this guide documents the port state as *blocking* instead of *discarding*. Designated ports start in the listening state.

Rapid Convergence

The RSTP provides for rapid recovery of connectivity following the failure of a switch, a switch port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links as follows:

- **Edge ports**—If you configure a port as an edge port on an RSTP switch by using the **spanning-tree portfast** interface configuration command, the edge port immediately transitions to the forwarding state. An edge port is the same as a Port Fast-enabled port, and you should enable it only on ports that connect to a single end station.
- **Root ports**—If the RSTP selects a new root port, it blocks the old root port and immediately transitions the new root port to the forwarding state.
- **Point-to-point links**—If you connect a port to another port through a point-to-point link and the local port becomes a designated port, it negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology.

As shown in Figure 15-2, Switch A is connected to Switch B through a point-to-point link, and all of the ports are in the blocking state. Assume that the priority of Switch A is a smaller numerical value than the priority of Switch B. Switch A sends a proposal message (a configuration BPDU with the proposal flag set) to Switch B, proposing itself as the designated switch.

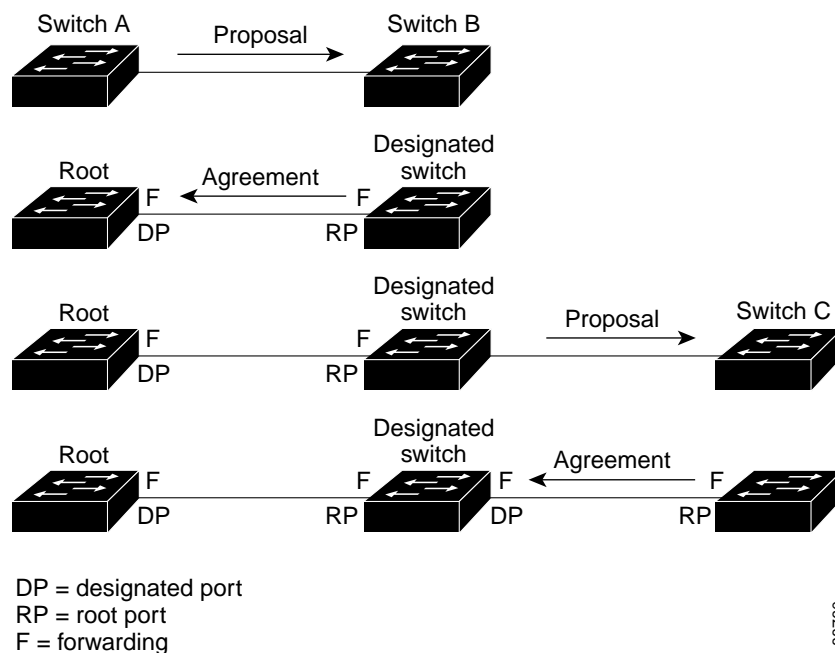
After receiving the proposal message, Switch B selects as its new root port the port from which the proposal message was received, forces all nonedge ports to the blocking state, and sends an agreement message (a BPDU with the agreement flag set) through its new root port.

After receiving Switch B's agreement message, Switch A also immediately transitions its designated port to the forwarding state. No loops in the network are formed because Switch B blocked all of its nonedge ports and because there is a point-to-point link between Switches A and B.

When Switch C is connected to Switch B, a similar set of handshaking messages are exchanged. Switch C selects the port connected to Switch B as its root port, and both ends immediately transition to the forwarding state. With each iteration of this handshaking process, one more switch joins the active topology. As the network converges, this proposal-agreement handshaking progresses from the root toward the leaves of the spanning tree.

The switch determines the link type from the port duplex mode: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. You can override the default setting that is determined by the duplex setting by using the **spanning-tree link-type** interface configuration command.

Figure 15-2 Proposal and Agreement Handshaking for Rapid Convergence



Synchronization of Port Roles

When the switch receives a proposal message on one of its ports and that port is selected as the new root port, the RSTP forces all other ports to synchronize with the new root information.

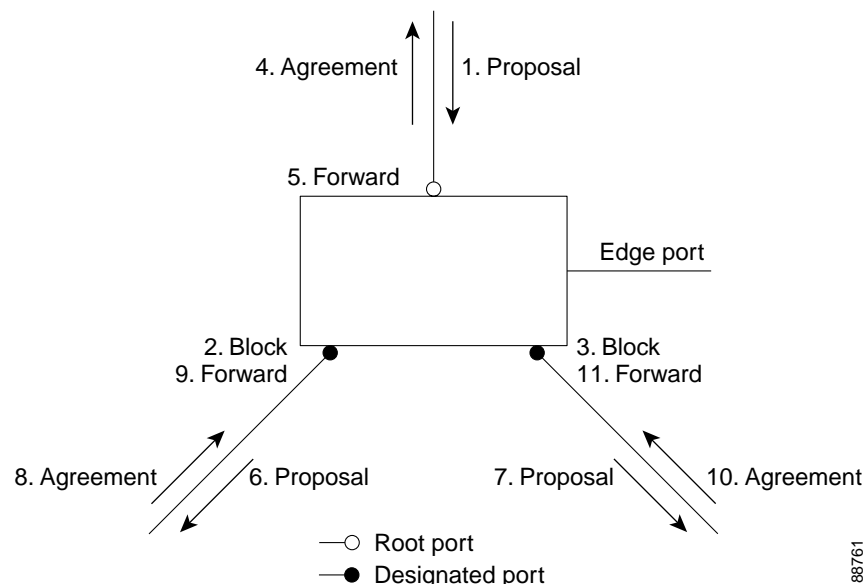
The switch is synchronized with superior root information received on the root port if all other ports are synchronized. An individual port on the switch is synchronized if

- That port is in the blocking state.
- It is an edge port (a port configured to be at the edge of the network).

If a designated port is in the forwarding state and is not configured as an edge port, it transitions to the blocking state when the RSTP forces it to synchronize with new root information. In general, when the RSTP forces a port to synchronize with root information and the port does not satisfy any of the above conditions, its port state is set to blocking.

After ensuring all of the ports are synchronized, the switch sends an agreement message to the designated switch corresponding to its root port. When the switches connected by a point-to-point link are in agreement about their port roles, the RSTP immediately transitions the port states to forwarding. The sequence of events is shown in [Figure 15-3](#).

Figure 15-3 Sequence of Events During Rapid Convergence



Bridge Protocol Data Unit Format and Processing

The RSTP BPDU format is the same as the IEEE 802.1D BPDU format except that the protocol version is set to 2. A new one-byte version 1 Length field is set to zero, which means that no version 1 protocol information is present. Table 15-2 shows the RSTP flag fields.

Table 15-2 RSTP BPDU Flags

Bit	Function
0	Topology change (TC)
1	Proposal
2–3:	Port role:
00	Unknown
01	Alternate port
10	Root port
11	Designated port
4	Learning
5	Forwarding
6	Agreement
7	Topology change acknowledgement (TCA)

The sending switch sets the proposal flag in the RSTP BPDU to propose itself as the designated switch on that LAN. The port role in the proposal message is always set to the designated port.

The sending switch sets the agreement flag in the RSTP BPDU to accept the previous proposal. The port role in the agreement message is always set to the root port.

The RSTP does not have a separate topology change notification (TCN) BPDU. It uses the topology change (TC) flag to show the topology changes. However, for interoperability with 802.1D switches, the RSTP switch processes and generates TCN BPDUs.

The learning and forwarding flags are set according to the state of the sending port.

Processing Superior BPDU Information

If a port receives superior root information (lower bridge ID, lower path cost, and so forth) than currently stored for the port, the RSTP triggers a reconfiguration. If the port is proposed and is selected as the new root port, RSTP forces all the other ports to synchronize.

If the BPDU received is an RSTP BPDU with the proposal flag set, the switch sends an agreement message after all of the other ports are synchronized. If the BPDU is an 802.1D BPDU, the switch does not set the proposal flag and starts the forward-delay timer for the port. The new root port requires twice the forward-delay time to transition to the forwarding state.

If the superior information received on the port causes the port to become a backup or alternate port, RSTP sets the port to the blocking state but does not send the agreement message. The designated port continues sending BPDUs with the proposal flag set until the forward-delay timer expires, at which time the port transitions to the forwarding state.

Processing Inferior BPDU Information

If a designated port receives an inferior BPDU (higher bridge ID, higher path cost, and so forth than currently stored for the port) with a designated port role, it immediately replies with its own information.

Topology Changes

This section describes the differences between the RSTP and the 802.1D in handling spanning-tree topology changes.

- **Detection**—Unlike 802.1D in which *any* transition between the blocking and the forwarding state causes a topology change, *only* transitions from the blocking to the forwarding state cause a topology change with RSTP (only an increase in connectivity is considered a topology change). State changes on an edge port do not cause a topology change. When an RSTP switch detects a topology change, it flushes the learned information on all of its nonedge ports except on those from which it received the TC notification.
- **Notification**—Unlike 802.1D, which uses TCN BPDUs, the RSTP does not use them. However, for 802.1D interoperability, an RSTP switch processes and generates TCN BPDUs.
- **Acknowledgement**—When an RSTP switch receives a TCN message on a designated port from an 802.1D switch, it replies with an 802.1D configuration BPDU with the TCA bit set. However, if the TC-while timer (the same as the topology-change timer in 802.1D) is active on a root port connected to an 802.1D switch and a configuration BPDU with the TCA bit set is received, the TC-while timer is reset.

This behavior is only required to support 802.1D switches. The RSTP BPDUs never have the TCA bit set.

- **Propagation**—When an RSTP switch receives a TC message from another switch through a designated or root port, it propagates the change to all of its nonedge, designated ports and to the root port (excluding the port on which it is received). The switch starts the TC-while timer for all such ports and flushes the information learned on them.
- **Protocol migration**—For backward compatibility with 802.1D switches, RSTP selectively sends 802.1D configuration BPDUs and TCN BPDUs on a per-port basis.

When a port is initialized, the migrate-delay timer is started (specifies the minimum time during which RSTP BPDUs are sent), and RSTP BPDUs are sent. While this timer is active, the switch processes all BPDUs received on that port and ignores the protocol type.

If the switch receives an 802.1D BPDU after the port's migration-delay timer has expired, it assumes that it is connected to an 802.1D switch and starts using only 802.1D BPDUs. However, if the RSTP switch is using 802.1D BPDUs on a port and receives an RSTP BPDU after the timer has expired, it restarts the timer and starts using RSTP BPDUs on that port.

Configuring MSTP Features

These sections describe how to configure basic MSTP features:

- [Default MSTP Configuration, page 15-12](#)
- [MSTP Configuration Guidelines, page 15-12](#)
- [Specifying the MST Region Configuration and Enabling MSTP, page 15-13](#) (required)
- [Configuring the Root Switch, page 15-14](#) (optional)
- [Configuring a Secondary Root Switch, page 15-16](#) (optional)
- [Configuring Port Priority, page 15-17](#) (optional)
- [Configuring Path Cost, page 15-18](#) (optional)
- [Configuring the Switch Priority, page 15-19](#) (optional)
- [Configuring the Hello Time, page 15-19](#) (optional)
- [Configuring the Forwarding-Delay Time, page 15-20](#) (optional)
- [Configuring the Maximum-Aging Time, page 15-21](#) (optional)
- [Configuring the Maximum-Hop Count, page 15-21](#) (optional)
- [Specifying the Link Type to Ensure Rapid Transitions, page 15-22](#) (optional)
- [Restarting the Protocol Migration Process, page 15-22](#) (optional)

Default MSTP Configuration

Table 15-3 shows the default MSTP configuration.

Table 15-3 Default MSTP Configuration

Feature	Default Setting
Spanning-tree mode	PVST+ (Rapid PVST+ and MSTP are disabled).
Switch priority (configurable on a per-CIST interface basis)	32768.
Spanning-tree port priority (configurable on a per-CIST interface basis)	128.
Spanning-tree port cost (configurable on a per-CIST interface basis)	1000 Mbps: 4. 100 Mbps: 19. 10 Mbps: 100.
Hello time	2 seconds.
Forward-delay time	15 seconds.
Maximum-aging time	20 seconds.
Maximum hop count	20 hops.

For information about the supported number of spanning-tree instances, see the [“Supported Spanning-Tree Instances”](#) section on page 14-9.

MSTP Configuration Guidelines

These are the configuration guidelines for MSTP:

- When you enable MST by using the **spanning-tree mode mst** global configuration command, RSTP is automatically enabled.
- For two or more switches to be in the same MST region, they must have the same VLAN-to-instance map, the same configuration revision number, and the same name.
- The switch supports up to 16 MST instances. The number of VLANs that can be mapped to a particular MST instance is unlimited.
- The UplinkFast and BackboneFast features are not supported with the MSTP.
- PVST+, rapid PVST+, and MSTP are supported, but only one version can be active at any time. (For example, all VLANs run PVST+, all VLANs run rapid PVST+, or all VLANs run MSTP.) For more information, see the [“Spanning-Tree Interoperability and Backward Compatibility”](#) section on page 14-10. For information on the recommended trunk port configuration, see the [“Interaction with Other Features”](#) section on page 11-18.
- VTP propagation of the MST configuration is not supported. However, you can manually configure the MST configuration (region name, revision number, and VLAN-to-instance mapping) on each switch within the MST region by using the command-line interface (CLI) or through the SNMP support.
- For load balancing across redundant paths in the network to work, all VLAN-to-instance mapping assignments must match; otherwise, all traffic flows on a single link.

- All MST boundary ports must be forwarding for load balancing between a PVST+ and an MST cloud or between a rapid-PVST+ and an MST cloud. For this to occur, the IST master of the MST cloud should also be the root of the CST. If the MST cloud consists of multiple MST regions, one of the MST regions must contain the CST root, and all of the other MST regions must have a better path to the root contained within the MST cloud than a path through the PVST+ or rapid-PVST+ cloud. You might have to manually configure the switches in the clouds.
- Partitioning the network into a large number of regions is not recommended. However, if this situation is unavoidable, we recommend that you partition the switched LAN into smaller LANs interconnected by routers or non-Layer 2 devices.


Specifying the MST Region Configuration and Enabling MSTP

For two or more switches to be in the same MST region, they must have the same VLAN-to-instance mapping, the same configuration revision number, and the same name.

A region can have one member or multiple members with the same MST configuration; each member must be capable of processing RSTP BPDUs. There is no limit to the number of MST regions in a network, but each region can support up to 16 spanning-tree instances. You can assign a VLAN to only one spanning-tree instance at a time.

Beginning in privileged EXEC mode, follow these steps to specify the MST region configuration and enable MSTP. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst configuration	Enter MST configuration mode.
Step 3	instance <i>instance-id</i> vlan <i>vlan-range</i>	Map VLANs to an MST instance. <ul style="list-style-type: none"> • For <i>instance-id</i>, the range is 1 to 15. • For vlan <i>vlan-range</i>, the range is 1 to 4094. <p>When you map VLANs to an MST instance, the mapping is incremental, and the VLANs specified in the command are added to or removed from the VLANs that were previously mapped.</p> <p>To specify a VLAN range, use a hyphen; for example, instance 1 vlan 1-63 maps VLANs 1 through 63 to MST instance 1.</p> <p>To specify a VLAN series, use a comma; for example, instance 1 vlan 10, 20, 30 maps VLANs 10, 20, and 30 to MST instance 1.</p>
Step 4	name <i>name</i>	Specify the configuration name. The <i>name</i> string has a maximum length of 32 characters and is case sensitive.
Step 5	revision <i>version</i>	Specify the configuration revision number. The range is 0 to 65535.
Step 6	show pending	Verify your configuration by displaying the pending configuration.
Step 7	exit	Apply all changes, and return to global configuration mode.

	Command	Purpose
Step 8	spanning-tree mode mst	Enable MSTP. RSTP is also enabled. <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> Caution Changing spanning-tree modes can disrupt traffic because all spanning-tree instances are stopped for the previous mode and restarted in the new mode. </div> </div> You cannot run both MSTP and PVST+ or both MSTP and rapid PVST+ at the same time.
Step 9	end	Return to privileged EXEC mode.
Step 10	show running-config	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default MST region configuration, use the **no spanning-tree mst configuration** global configuration command. To return to the default VLAN-to-instance map, use the **no instance *instance-id* [vlan *vlan-range*]** MST configuration command. To return to the default name, use the **no name** MST configuration command. To return to the default revision number, use the **no revision** MST configuration command. To re-enable PVST+, use the **no spanning-tree mode** or the **spanning-tree mode pvst** global configuration command.

This example shows how to enter MST configuration mode, map VLANs 10 to 20 to MST instance 1, name the region *region1*, set the configuration revision to 1, display the pending configuration, apply the changes, and return to global configuration mode:

```
Switch(config)# spanning-tree mst configuration
Switch(config-mst)# instance 1 vlan 10-20
Switch(config-mst)# name region1
Switch(config-mst)# revision 1
Switch(config-mst)# show pending
Pending MST configuration
Name      [region1]
Revision  1
Instance  Vlan  Mapped
-----  -
0         1-9, 21-4094
1         10-20
-----

Switch(config-mst)# exit
Switch(config)#
```

Configuring the Root Switch

The switch maintains a spanning-tree instance for the group of VLANs mapped to it. A bridge ID, consisting of the switch priority and the switch MAC address, is associated with each instance. For a group of VLANs, the switch with the lowest bridge ID becomes the root switch.

To configure a switch to become the root, use the **spanning-tree mst *instance-id* root** global configuration command to modify the switch priority from the default value (32768) to a significantly lower value so that the switch becomes the root switch for the specified spanning-tree instance. When you enter this command, the switch checks the switch priorities of the root switches. Because of the extended system ID support, the switch sets its own priority for the specified instance to 24576 if this value will cause this switch to become the root for the specified spanning-tree instance.

If any root switch for the specified instance has a switch priority lower than 24576, the switch sets its own priority to 4096 less than the lowest switch priority. (4096 is the value of the least-significant bit of a 4-bit switch priority value as shown in [Table 14-1 on page 14-4](#).)

**Note**

Catalyst 2970 switches running software earlier than Cisco IOS Release 12.1(14)EA1 do not support the MSTP.

**Note**

If your network consists of switches that both do and do not support the extended system ID, it is unlikely that the switch with the extended system ID support will become the root switch. The extended system ID increases the switch priority value every time the VLAN number is greater than the priority of the connected switches running older software.

**Note**

The root switch for each spanning-tree instance should be a backbone or distribution switch. Do not configure an access switch as the spanning-tree primary root.

Use the **diameter** keyword, which is available only for MST instance 0, to specify the Layer 2 network diameter (that is, the maximum number of switch hops between any two end stations in the Layer 2 network). When you specify the network diameter, the switch automatically sets an optimal hello time, forward-delay time, and maximum-age time for a network of that diameter, which can significantly reduce the convergence time. You can use the **hello** keyword to override the automatically calculated hello time.

**Note**

After configuring the switch as the root switch, we recommend that you avoid manually configuring the hello time, forward-delay time, and maximum-age time through the **spanning-tree mst hello-time**, **spanning-tree mst forward-time**, and the **spanning-tree mst max-age** global configuration commands.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the root switch. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst <i>instance-id</i> root primary [<i>diameter net-diameter</i> [<i>hello-time seconds</i>]]	Configure a switch as the root switch. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. (Optional) For diameter net-diameter, specify the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. (Optional) For hello-time seconds, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds.

	Command	Purpose
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst *instance-id* root** global configuration command.

Configuring a Secondary Root Switch

When you configure a Catalyst 2970 switch with the extended system ID support as the secondary root, the switch priority is modified from the default value (32768) to 28672. The switch is then likely to become the root switch for the specified instance if the primary root switch fails. This is assuming that the other network switches use the default switch priority of 32768 and therefore are unlikely to become the root switch.

You can execute this command on more than one switch to configure multiple backup root switches. Use the same network diameter and hello-time values that you used when you configured the primary root switch with the **spanning-tree mst *instance-id* root primary** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure a switch as the secondary root switch. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst <i>instance-id</i> root secondary [diameter <i>net-diameter</i> [hello-time <i>seconds</i>]]	Configure a switch as the secondary root switch. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. (Optional) For diameter <i>net-diameter</i>, specify the maximum number of switches between any two end stations. The range is 2 to 7. This keyword is available only for MST instance 0. (Optional) For hello-time <i>seconds</i>, specify the interval in seconds between the generation of configuration messages by the root switch. The range is 1 to 10 seconds; the default is 2 seconds. Use the same network diameter and hello-time values that you used when configuring the primary root switch. See the “Configuring the Root Switch” section on page 15-14 .
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst *instance-id* root** global configuration command.

Configuring Port Priority

If a loop occurs, the MSTP uses the port priority when selecting an interface to put into the forwarding state. You can assign higher priority values (lower numerical values) to interfaces that you want selected first and lower priority values (higher numerical values) that you want selected last. If all interfaces have the same priority value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the MSTP port priority of an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces. The port-channel range is 1 to 12.
Step 3	spanning-tree mst <i>instance-id</i> port-priority <i>priority</i>	Configure the port priority. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>priority</i>, the range is 0 to 240 in increments of 16. The default is 128. The lower the number, the higher the priority. The priority values are 0, 16, 32, 48, 64, 80, 96, 112, 128, 144, 160, 176, 192, 208, 224, and 240. All other values are rejected.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

The **show spanning-tree mst interface** *interface-id* privileged EXEC command displays information only if the port is in a link-up operative state. Otherwise, you can use the **show running-config interface** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree mst** *instance-id* **port-priority** interface configuration command.

Configuring Path Cost

The MSTP path cost default value is derived from the media speed of an interface. If a loop occurs, the MSTP uses cost when selecting an interface to put in the forwarding state. You can assign lower cost values to interfaces that you want selected first and higher cost values that you want selected last. If all interfaces have the same cost value, the MSTP puts the interface with the lowest interface number in the forwarding state and blocks the other interfaces.

Beginning in privileged EXEC mode, follow these steps to configure the MSTP cost of an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure. Valid interfaces include physical interfaces and port-channel logical interfaces. The port-channel range is 1 to 12.
Step 3	spanning-tree mst <i>instance-id</i> cost <i>cost</i>	Configure the cost. If a loop occurs, the MSTP uses the path cost when selecting an interface to place into the forwarding state. A lower path cost represents higher-speed transmission. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>cost</i>, the range is 1 to 200000000; the default value is derived from the media speed of the interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree mst interface <i>interface-id</i> or show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

The **show spanning-tree mst interface** *interface-id* privileged EXEC command displays information only for ports that are in a link-up operative state. Otherwise, you can use the **show running-config** privileged EXEC command to confirm the configuration.

To return the interface to its default setting, use the **no spanning-tree mst** *instance-id* **cost** interface configuration command.

Configuring the Switch Priority

You can configure the switch priority and make it more likely that the switch will be chosen as the root switch.



Note

Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree mst *instance-id* root primary** and the **spanning-tree mst *instance-id* root secondary** global configuration commands to modify the switch priority.

Beginning in privileged EXEC mode, follow these steps to configure the switch priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst <i>instance-id</i> priority <i>priority</i>	Configure the switch priority. <ul style="list-style-type: none"> For <i>instance-id</i>, you can specify a single instance, a range of instances separated by a hyphen, or a series of instances separated by a comma. The range is 0 to 15. For <i>priority</i>, the range is 0 to 61440 in increments of 4096; the default is 32768. The lower the number, the more likely the switch will be chosen as the root switch. Priority values are 0, 4096, 8192, 12288, 16384, 20480, 24576, 28672, 32768, 36864, 40960, 45056, 49152, 53248, 57344, and 61440. All other values are rejected.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst <i>instance-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst *instance-id* priority** global configuration command.

Configuring the Hello Time

You can configure the interval between the generation of configuration messages by the root switch by changing the hello time.



Note

Exercise care when using this command. For most situations, we recommend that you use the **spanning-tree mst *instance-id* root primary** and the **spanning-tree mst *instance-id* root secondary** global configuration commands to modify the hello time.

Beginning in privileged EXEC mode, follow these steps to configure the hello time for all MST instances. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst hello-time <i>seconds</i>	Configure the hello time for all MST instances. The hello time is the interval between the generation of configuration messages by the root switch. These messages mean that the switch is alive. For <i>seconds</i> , the range is 1 to 10; the default is 2.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst hello-time** global configuration command.

Configuring the Forwarding-Delay Time

Beginning in privileged EXEC mode, follow these steps to configure the forwarding-delay time for all MST instances. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst forward-time <i>seconds</i>	Configure the forward time for all MST instances. The forward delay is the number of seconds a port waits before changing from its spanning-tree learning and listening states to the forwarding state. For <i>seconds</i> , the range is 4 to 30; the default is 15.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst forward-time** global configuration command.

Configuring the Maximum-Aging Time

Beginning in privileged EXEC mode, follow these steps to configure the maximum-aging time for all MST instances. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst max-age <i>seconds</i>	Configure the maximum-aging time for all MST instances. The maximum-aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration. For <i>seconds</i> , the range is 6 to 40; the default is 20.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst max-age** global configuration command.

Configuring the Maximum-Hop Count

Beginning in privileged EXEC mode, follow these steps to configure the maximum-hop count for all MST instances. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree mst max-hops <i>hop-count</i>	Specify the number of hops in a region before the BPDU is discarded, and the information held for a port is aged. For <i>hop-count</i> , the range is 1 to 40; the default is 20.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default setting, use the **no spanning-tree mst max-hops** global configuration command.

Specifying the Link Type to Ensure Rapid Transitions

If you connect a port to another port through a point-to-point link and the local port becomes a designated port, the RSTP negotiates a rapid transition with the other port by using the proposal-agreement handshake to ensure a loop-free topology as described in the [“Rapid Convergence” section on page 15-7](#).

By default, the link type is determined from the duplex mode of the interface: a full-duplex port is considered to have a point-to-point connection; a half-duplex port is considered to have a shared connection. If you have a half-duplex link physically connected point-to-point to a single port on a remote switch running MSTP, you can override the default setting of the link type and enable rapid transitions to the forwarding state.

Beginning in privileged EXEC mode, follow these steps to override the default link-type setting. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 1	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure. Valid interfaces include physical interface, VLANs, and port-channel logical interfaces. The VLAN ID range is 1 to 4094. The port-channel range is 1 to 12.
Step 2	spanning-tree link-type point-to-point	Specify that the link type of a port is point-to-point.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree mst interface <i>interface-id</i>	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the port to its default setting, use the **no spanning-tree link-type** interface configuration command.

Restarting the Protocol Migration Process

A switch running MSTP supports a built-in protocol migration mechanism that enables it to interoperate with legacy 802.1D switches. If this switch receives a legacy 802.1D configuration BPDU (a BPDU with the protocol version set to 0), it sends only 802.1D BPDUs on that port. An MSTP switch also can detect that a port is at the boundary of a region when it receives a legacy BPDU, an MST BPDU (version 3) associated with a different region, or an RST BPDU (version 2).

However, the switch does not automatically revert to the MSTP mode if it no longer receives 802.1D BPDUs because it cannot determine whether the legacy switch has been removed from the link unless the legacy switch is the designated switch. A switch also might continue to assign a boundary role to a port when the switch to which it is connected has joined the region.

To restart the protocol migration process (force the renegotiation with neighboring switches) on the switch, use the **clear spanning-tree detected-protocols** privileged EXEC command.

To restart the protocol migration process on a specific interface, use the **clear spanning-tree detected-protocols interface *interface-id*** privileged EXEC command.

Displaying the MST Configuration and Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in [Table 15-4](#):

Table 15-4 *Commands for Displaying MST Status*

Command	Purpose
show spanning-tree mst configuration	Displays the MST region configuration.
show spanning-tree mst <i>instance-id</i>	Displays MST information for the specified instance.
show spanning-tree mst interface <i>interface-id</i>	Displays MST information for the specified interface.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.



Configuring Optional Spanning-Tree Features

This chapter describes how to configure optional spanning-tree features on the Catalyst 2970 switch. You can configure all of these features when your switch is running the per-VLAN spanning-tree plus (PVST+). You can configure only the noted features when your switch is running the Multiple Spanning Tree Protocol (MSTP) or the rapid per-VLAN spanning-tree plus (rapid-PVST+) protocol.

For information on configuring the PVST+ and rapid PVST+, see [Chapter 14, “Configuring STP.”](#) For information about the Multiple Spanning Tree Protocol (MSTP) and how to map multiple VLANs to the same spanning-tree instance, see [Chapter 15, “Configuring MSTP.”](#)



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding Optional Spanning-Tree Features, page 16-1](#)
- [Configuring Optional Spanning-Tree Features, page 16-9](#)
- [Displaying the Spanning-Tree Status, page 16-16](#)

Understanding Optional Spanning-Tree Features

These sections describe how the optional spanning-tree features work:

- [Understanding Port Fast, page 16-2](#)
- [Understanding BPDU Guard, page 16-3](#)
- [Understanding BPDU Filtering, page 16-3](#)
- [Understanding UplinkFast, page 16-4](#)
- [Understanding BackboneFast, page 16-5](#)
- [Understanding Root Guard, page 16-7](#)
- [Understanding Loop Guard, page 16-8](#)

Understanding Port Fast

Port Fast immediately brings an interface configured as an access or trunk port to the forwarding state from a blocking state, bypassing the listening and learning states. You can use Port Fast on ports connected to a single workstation or server, as shown in [Figure 16-1](#), to allow those devices to immediately connect to the network, rather than waiting for the spanning tree to converge.

Ports connected to a single workstation or server should not receive bridge protocol data units (BPDUs). A port with Port Fast enabled goes through the normal cycle of spanning-tree status changes when the switch is restarted.

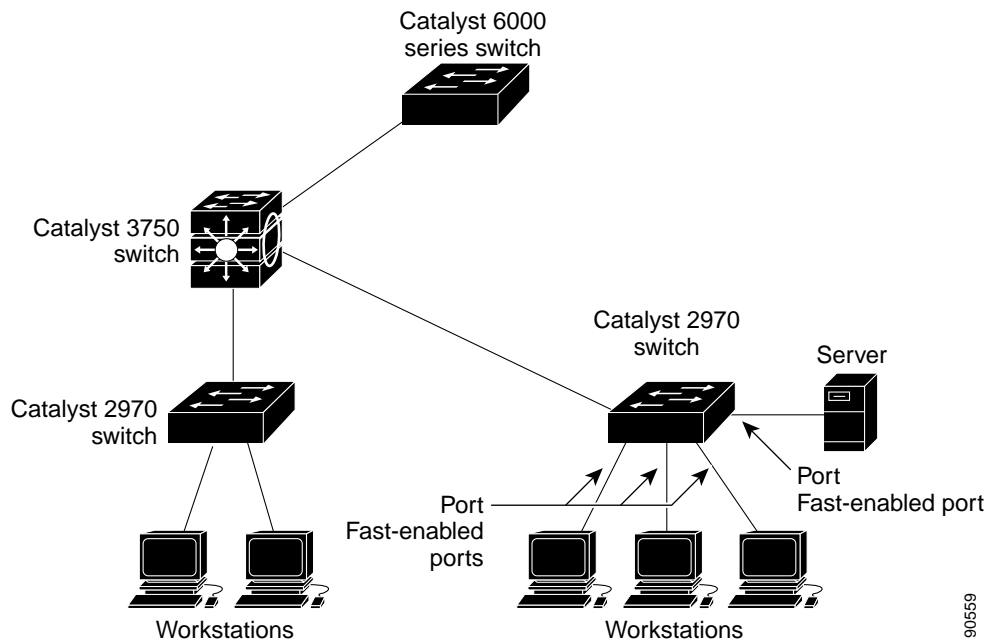


Note

Because the purpose of Port Fast is to minimize the time ports must wait for spanning-tree to converge, it is effective only when used on ports connected to end stations. If you enable Port Fast on a port connecting to another switch, you risk creating a spanning-tree loop.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree portfast** interface configuration or the **spanning-tree portfast default** global configuration command.

Figure 16-1 Port Fast-Enabled Ports



Understanding BPDU Guard

The BPDU guard feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you enable BPDU guard on Port Fast-enabled ports by using the **spanning-tree portfast bpduguard default** global configuration command. Spanning tree shuts down ports that are in a Port Fast-operational state. In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state.

At the interface level, you enable BPDU guard on any port by using the **spanning-tree bpduguard enable** interface configuration command without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

The BPDU guard feature provides a secure response to invalid configurations because you must manually put the port back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable the BPDU guard feature for the entire switch or for an interface.

Understanding BPDU Filtering

The BPDU filtering feature can be globally enabled on the switch or can be enabled per interface, but the feature operates with some differences.

At the global level, you can enable BPDU filtering on Port Fast-enabled ports by using the **spanning-tree portfast bpdupfilter default** global configuration command. This command prevents ports that are in a Port Fast-operational state from sending or receiving BPDUs. The ports still send a few BPDUs at link-up before the switch begins to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a Port Fast-enabled port, the port loses its Port Fast-operational status, and BPDU filtering is disabled.

At the interface level, you can enable BPDU filtering on any port by using the **spanning-tree bpdupfilter enable** interface configuration command without also enabling the Port Fast feature. This command prevents the port from sending or receiving BPDUs.



Caution

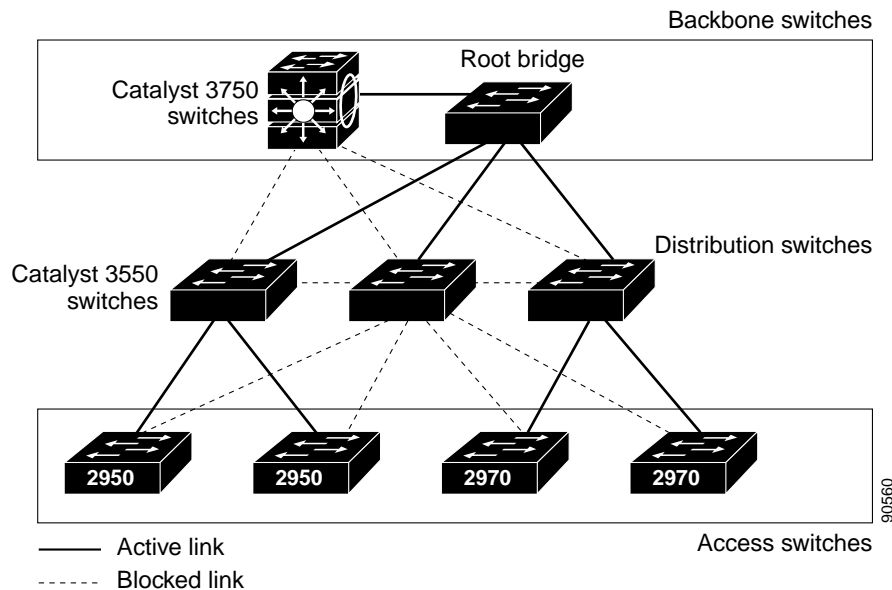
Enabling BPDU filtering on an interface is the same as disabling spanning tree on it and can result in spanning-tree loops.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable the BPDU filtering feature for the entire switch or for an interface.

Understanding UplinkFast

Switches in hierarchical networks can be grouped into backbone switches, distribution switches, and access switches. Figure 16-2 shows a complex network where distribution switches and access switches each have at least one redundant link that spanning tree blocks to prevent loops.

Figure 16-2 Switches in a Hierarchical Network



If a switch loses connectivity, it begins using the alternate paths as soon as the spanning tree selects a new root port. By enabling UplinkFast with the **spanning-tree uplinkfast** global configuration command, you can accelerate the choice of a new root port when a link or switch fails or when the spanning tree reconfigures itself. The root port transitions to the forwarding state immediately without going through the listening and learning states, as it would with the normal spanning-tree procedures. The UplinkFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP because these protocols use fast convergence and take precedence over UplinkFast.

When the spanning tree reconfigures the new root port, other interfaces flood the network with multicast packets, one for each address that was learned on the interface. You can limit these bursts of multicast traffic by reducing the max-update-rate parameter (the default for this parameter is 150 packets per second). However, if you enter zero, station-learning frames are not generated, so the spanning-tree topology converges more slowly after a loss of connectivity.



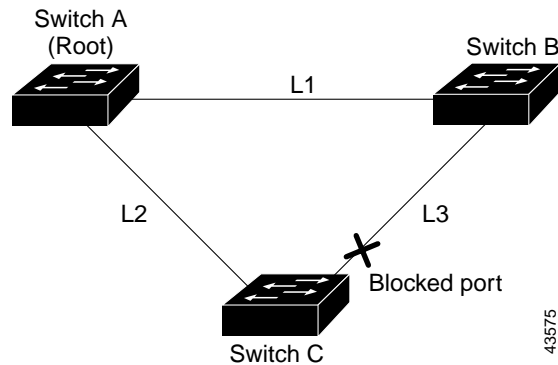
Note

UplinkFast is most useful in wiring-closet switches at the access or edge of the network. It is not appropriate for backbone devices. This feature might not be useful for other types of applications.

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links using uplink groups. An uplink group is a set of Layer 2 interfaces (per VLAN), only one of which is forwarding at any given time. Specifically, an uplink group consists of the root port (which is forwarding) and a set of blocked ports, except for self-looping ports. The uplink group provides an alternate path in case the currently forwarding link fails.

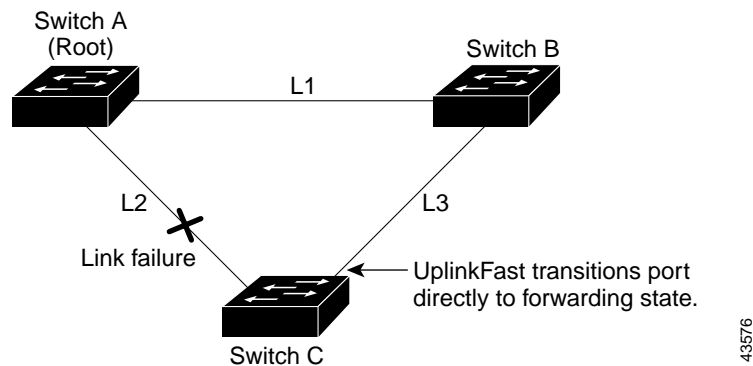
Figure 16-3 shows an example topology with no link failures. Switch A, the root switch, is connected directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that is connected directly to Switch B is in a blocking state.

Figure 16-3 UplinkFast Example Before Direct Link Failure



If Switch C detects a link failure on the currently active link L2 on the root port (a *direct* link failure), UplinkFast unblocks the blocked port on Switch C and transitions it to the forwarding state without going through the listening and learning states, as shown in Figure 16-4. This change takes approximately 1 to 5 seconds.

Figure 16-4 UplinkFast Example After Direct Link Failure



Understanding BackboneFast

BackboneFast detects indirect failures in the core of the backbone. BackboneFast is a complementary technology to the UplinkFast feature, which responds to failures on links directly connected to access switches. BackboneFast optimizes the maximum-age timer, which determines the amount of time the switch stores protocol information received on an interface. When a switch receives an inferior BPDU from the designated port of another switch, the BPDU is a signal that the other switch might have lost its path to the root, and BackboneFast tries to find an alternate path to the root. The BackboneFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

BackboneFast, which is enabled by using the **spanning-tree backbonefast** global configuration command, starts when a root port or blocked port on a switch receives inferior BPDUs from its designated switch. An inferior BPDU identifies a switch that declares itself as both the root bridge and

the designated switch. When a switch receives an inferior BPDU, it means that a link to which the switch is not directly connected (an *indirect* link) has failed (that is, the designated switch has lost its connection to the root switch). Under spanning-tree rules, the switch ignores inferior BPDUs for the configured maximum aging time specified by the **spanning-tree vlan *vlan-id* max-age** global configuration command.

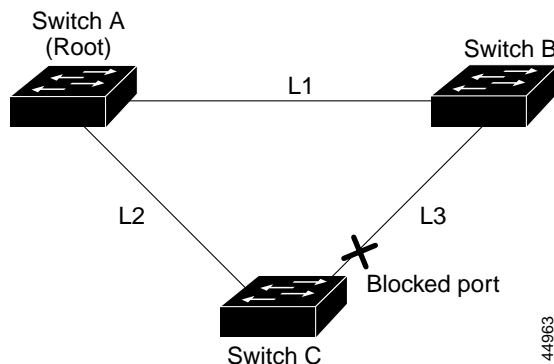
The switch tries to determine if it has an alternate path to the root switch. If the inferior BPDU arrives on a blocked port, the root port and other blocked ports on the switch become alternate paths to the root switch. (Self-looped ports are not considered alternate paths to the root switch.) If the inferior BPDU arrives on the root port, all blocked ports become alternate paths to the root switch. If the inferior BPDU arrives on the root port and there are no blocked ports, the switch assumes that it has lost connectivity to the root switch, causes the maximum aging time on the root port to expire, and becomes the root switch according to normal spanning-tree rules.

If the switch has alternate paths to the root switch, it uses these alternate paths to send a root link query (RLQ) request. The switch sends the RLQ request on all alternate paths and waits for an RLQ reply from other switches in the network.

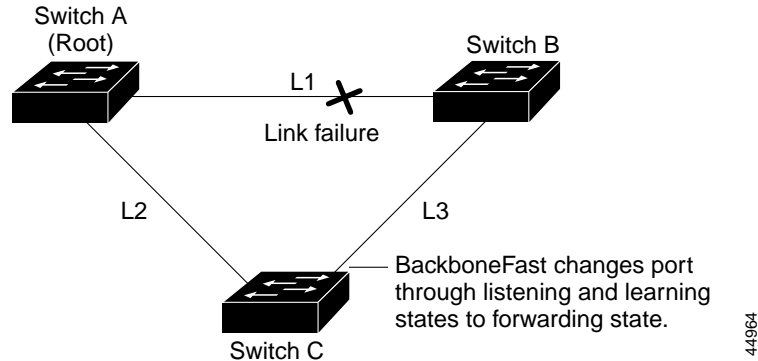
If the switch determines that it still has an alternate path to the root, it expires the maximum aging time on the port that received the inferior BPDU. If all the alternate paths to the root switch indicate that the switch has lost connectivity to the root switch, the switch expires the maximum aging time on the port that received the RLQ reply. If one or more alternate paths can still connect to the root switch, the switch makes all ports on which it received an inferior BPDU its designated ports and moves them from the blocking state (if they were in the blocking state), through the listening and learning states, and into the forwarding state.

Figure 16-5 shows an example topology with no link failures. Switch A, the root switch, connects directly to Switch B over link L1 and to Switch C over link L2. The Layer 2 interface on Switch C that connects directly to Switch B is in the blocking state.

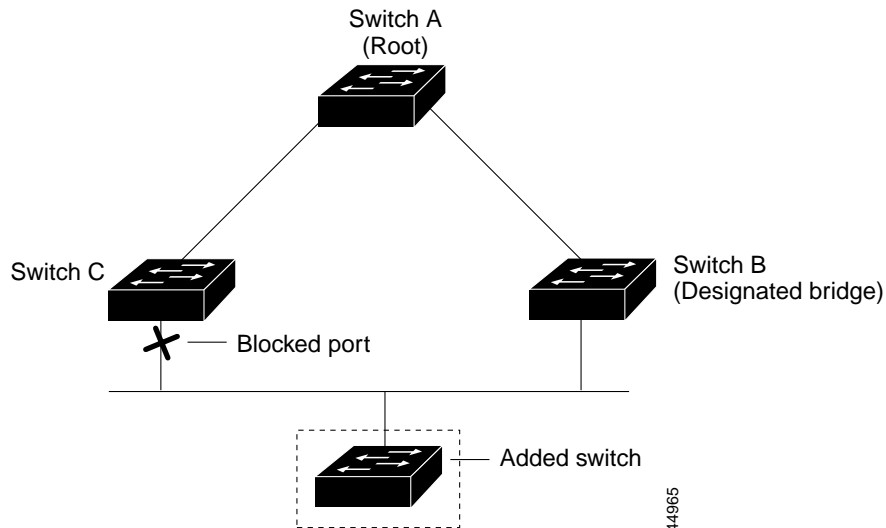
Figure 16-5 BackboneFast Example Before Indirect Link Failure



If link L1 fails as shown in Figure 16-6, Switch C cannot detect this failure because it is not connected directly to link L1. However, because Switch B is directly connected to the root switch over L1, it detects the failure, elects itself the root, and begins sending BPDUs to Switch C, identifying itself as the root. When Switch C receives the inferior BPDUs from Switch B, Switch C assumes that an indirect failure has occurred. At that point, BackboneFast allows the blocked port on Switch C to move immediately to the listening state without waiting for the maximum aging time for the port to expire. BackboneFast then transitions the Layer 2 interface on Switch C to the forwarding state, providing a path from Switch B to Switch A. This switchover takes approximately 30 seconds, twice the Forward Delay time if the default Forward Delay time of 15 seconds is set. Figure 16-6 shows how BackboneFast reconfigures the topology to account for the failure of link L1.

Figure 16-6 BackboneFast Example After Indirect Link Failure

If a new switch is introduced into a shared-medium topology as shown in [Figure 16-7](#), BackboneFast is not activated because the inferior BPDUs did not come from the recognized designated switch (Switch B). The new switch begins sending inferior BPDUs that indicate it is the root switch. However, the other switches ignore these inferior BPDUs, and the new switch learns that Switch B is the designated switch to Switch A, the root switch.

Figure 16-7 Adding a Switch in a Shared-Medium Topology

Understanding Root Guard

The Layer 2 network of a service provider (SP) can include many connections to switches that are not owned by the SP. In such a topology, the spanning tree can reconfigure itself and select a *customer switch* as the root switch, as shown in [Figure 16-8](#). You can avoid this situation by enabling root guard on SP switch interfaces that connect to switches in your customer's network. If spanning-tree calculations cause an interface in the customer network to be selected as the root port, root guard then places the interface in the root-inconsistent (blocked) state to prevent the customer's switch from becoming the root switch or being in the path to the root.

If a switch outside the SP network becomes the root switch, the interface is blocked (root-inconsistent state), and spanning tree selects a new root switch. The customer's switch does not become the root switch and is not in the path to the root.

If the switch is operating in multiple spanning-tree (MST) mode, root guard forces the port to be a designated port. If a boundary port is blocked in an internal spanning-tree (IST) instance because of root guard, the port also is blocked in all MST instances. A boundary port is a port that connects to a LAN, the designated switch of which is either an 802.1D switch or a switch with a different MST region configuration.

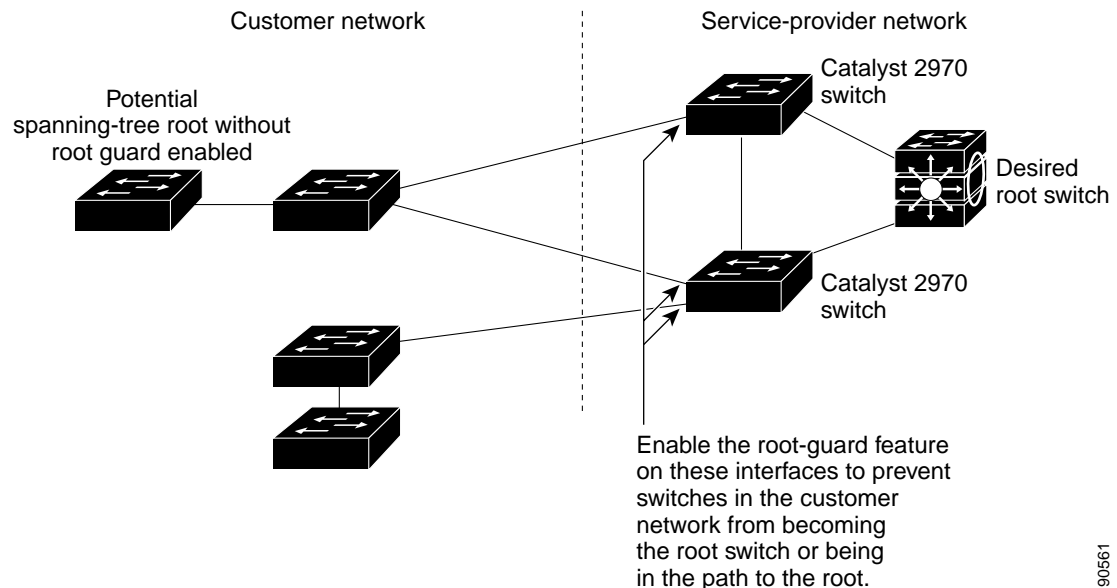
Root guard enabled on an interface applies to all the VLANs to which the interface belongs. VLANs can be grouped and mapped to an MST instance.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree guard root** interface configuration command.



Misuse of the root-guard feature can cause a loss of connectivity.

Figure 16-8 Root Guard in a Service-Provider Network



90561

Understanding Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is enabled on the entire switched network. Loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

If your switch is running PVST+, rapid PVST+, or MSTP, you can enable this feature by using the **spanning-tree loopguard default** global configuration command.

When the switch is operating in PVST+ or rapid-PVST+ mode, loop guard prevents alternate and root ports from becoming designated ports, and spanning tree does not send BPDUs on root or alternate ports.

When the switch is operating in MST mode, BPDUs are not sent on nonboundary ports only if the port is blocked by loop guard in all MST instances. On a boundary port, loop guard blocks the port in all MST instances.

Configuring Optional Spanning-Tree Features

These sections describe how to configure optional spanning-tree features:

- [Default Optional Spanning-Tree Configuration, page 16-9](#)
- [Optional Spanning-Tree Configuration Guidelines, page 16-9](#)
- [Enabling Port Fast, page 16-10](#) (optional)
- [Enabling BPDU Guard, page 16-11](#) (optional)
- [Enabling BPDU Filtering, page 16-12](#) (optional)
- [Enabling UplinkFast for Use with Redundant Links, page 16-13](#) (optional)
- [Enabling BackboneFast, page 16-13](#) (optional)
- [Enabling Root Guard, page 16-14](#) (optional)
- [Enabling Loop Guard, page 16-15](#) (optional)

Default Optional Spanning-Tree Configuration

[Table 16-1](#) shows the default optional spanning-tree configuration.

Table 16-1 *Default Optional Spanning-Tree Configuration*

Feature	Default Setting
Port Fast, BPDU filtering, BPDU guard	Globally disabled (unless they are individually configured per interface).
UplinkFast	Globally disabled.
BackboneFast	Globally disabled.
Root guard	Disabled on all interfaces.
Loop guard	Disabled on all interfaces.

Optional Spanning-Tree Configuration Guidelines

The UplinkFast and BackboneFast features are not supported with the rapid PVST+ or the MSTP.

Enabling Port Fast

A port with the Port Fast feature enabled is moved directly to the spanning-tree forwarding state without waiting for the standard forward-time delay.




Caution

Use Port Fast *only* when connecting a single end station to an access or trunk port. Enabling this feature on a port connected to a switch or hub could prevent spanning tree from detecting and disabling loops in your network, which could cause broadcast storms and address-learning problems.

If you enable the voice VLAN feature, the Port Fast feature is automatically enabled. When you disable voice VLAN, the Port Fast feature is not automatically disabled. For more information, see [Chapter 13, “Configuring Voice VLAN.”](#)

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable Port Fast. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure.
Step 3	spanning-tree portfast [trunk]	<p>Enable Port Fast on an access port connected to a single workstation or server. By specifying the trunk keyword, you can enable Port Fast on a trunk port.</p> <div>  <p>Caution Make sure that there are no loops in the network between the trunk port and the workstation or server before you enable Port Fast on a trunk port.</p> </div> <p>By default, Port Fast is disabled on all ports.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show spanning-tree interface <i>interface-id</i> portfast	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

You can use the **spanning-tree portfast default** global configuration command to globally enable the Port Fast feature on all nontrunking ports.

To disable the Port Fast feature, use the **spanning-tree portfast disable** interface configuration command.

Enabling BPDU Guard

When you globally enable BPDU guard on ports that are Port Fast-enabled (the ports are in a Port Fast-operational state), spanning tree shuts down Port Fast-enabled ports that receive BPDUs.

In a valid configuration, Port Fast-enabled ports do not receive BPDUs. Receiving a BPDU on a Port Fast-enabled port signals an invalid configuration, such as the connection of an unauthorized device, and the BPDU guard feature puts the port in the error-disabled state. The BPDU guard feature provides a secure response to invalid configurations because you must manually put the port back in service. Use the BPDU guard feature in a service-provider network to prevent an access port from participating in the spanning tree.



Caution

Configure Port Fast only on ports that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You also can use the **spanning-tree bpduguard enable** interface configuration command to enable BPDU guard on any port without also enabling the Port Fast feature. When the port receives a BPDU, it is put in the error-disabled state.

You can enable the BPDU guard feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to globally enable the BPDU guard feature. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree portfast bpduguard default	Globally enable BPDU guard. By default, BPDU guard is disabled.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to an end station.
Step 4	spanning-tree portfast	Enable the Port Fast feature.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BPDU guard, use the **no spanning-tree portfast bpduguard default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpduguard default** global configuration command by using the **spanning-tree bpduguard enable** interface configuration command.

Enabling BPDU Filtering

When you globally enable BPDU filtering on Port Fast-enabled ports, it prevents ports that are in a Port Fast-operational state from sending or receiving BPDUs. The ports still send a few BPDUs at link-up before the switch begins to filter outbound BPDUs. You should globally enable BPDU filtering on a switch so that hosts connected to these ports do not receive BPDUs. If a BPDU is received on a Port Fast-enabled port, the port loses its Port Fast-operational status, and BPDU filtering is disabled.



Caution

Configure Port Fast only on ports that connect to end stations; otherwise, an accidental topology loop could cause a data packet loop and disrupt switch and network operation.

You can also use the **spanning-tree bpdupfilter enable** interface configuration command to enable BPDU filtering on any port without also enabling the Port Fast feature. This command prevents the port from sending or receiving BPDUs.



Caution

Enabling BPDU filtering on an interface is the same as disabling spanning tree on it and can result in spanning-tree loops.

You can enable the BPDU filtering feature if your switch is running PVST+, rapid PVST+, or MSTP. Beginning in privileged EXEC mode, follow these steps to globally enable the BPDU filtering feature. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree portfast bpdupfilter default	Globally enable BPDU filtering. By default, BPDU filtering is disabled.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to an end station.
Step 4	spanning-tree portfast	Enable the Port Fast feature.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BPDU filtering, use the **no spanning-tree portfast bpdupfilter default** global configuration command.

You can override the setting of the **no spanning-tree portfast bpdupfilter default** global configuration command by using the **spanning-tree bpdupfilter enable** interface configuration command.

Enabling UplinkFast for Use with Redundant Links

UplinkFast cannot be enabled on VLANs that have been configured with a switch priority. To enable UplinkFast on a VLAN with switch priority configured, first restore the switch priority on the VLAN to the default value by using the **no spanning-tree vlan *vlan-id* priority** global configuration command.



Note

When you enable UplinkFast, it affects all VLANs on the switch. You cannot configure UplinkFast on an individual VLAN.

The UplinkFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable UplinkFast. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree uplinkfast [max-update-rate <i>pkts-per-second</i>]	Enable UplinkFast. (Optional) For <i>pkts-per-second</i> , the range is 0 to 32000 packets per second; the default is 150. If you set the rate to 0, station-learning frames are not generated, and the spanning-tree topology converges more slowly after a loss of connectivity.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree summary	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

When UplinkFast is enabled, the switch priority of all VLANs is set to 49152. If you change the path cost to a value less than 3000 and you enable UplinkFast or UplinkFast is already enabled, the path cost of all interfaces and VLAN trunks is increased by 3000 (if you change the path cost to 3000 or above, the path cost is not altered). The changes to the switch priority and the path cost reduce the chance that a switch will become the root switch.

When UplinkFast is disabled, the switch priorities of all VLANs and path costs of all interfaces are set to default values if you did not modify them from their defaults.

To return the update packet rate to the default setting, use the **no spanning-tree uplinkfast max-update-rate** global configuration command. To disable UplinkFast, use the **no spanning-tree uplinkfast** command.

Enabling BackboneFast

You can enable BackboneFast to detect indirect link failures and to start the spanning-tree reconfiguration sooner.



Note

If you use BackboneFast, you must enable it on all switches in the network. BackboneFast is not supported on Token Ring VLANs. This feature is supported for use with third-party switches.

The BackboneFast feature is supported only when the switch is running PVST+. It is not supported when the switch is running rapid PVST+ or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable BackboneFast. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	spanning-tree backbonefast	Enable BackboneFast.
Step 3	end	Return to privileged EXEC mode.
Step 4	show spanning-tree summary	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the BackboneFast feature, use the **no spanning-tree backbonefast** global configuration command.

Enabling Root Guard

Root guard enabled on an interface applies to all the VLANs to which the interface belongs. Do not enable the root guard on interfaces to be used by the UplinkFast feature. With UplinkFast, the backup interfaces (in the blocked state) replace the root port in the case of a failure. However, if root guard is also enabled, all the backup interfaces used by the UplinkFast feature are placed in the root-inconsistent state (blocked) and are prevented from reaching the forwarding state.



Note

You cannot enable both root guard and loop guard at the same time.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable root guard on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify an interface to configure.
Step 3	spanning-tree guard root	Enable root guard on the interface. By default, root guard is disabled on all interfaces.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable root guard, use the **no spanning-tree guard** interface configuration command.

Enabling Loop Guard

You can use loop guard to prevent alternate or root ports from becoming designated ports because of a failure that leads to a unidirectional link. This feature is most effective when it is configured on the entire switched network. Loop guard operates only on ports that are considered point-to-point by the spanning tree.



Note

You cannot enable both loop guard and root guard at the same time.

You can enable this feature if your switch is running PVST+, rapid PVST+, or MSTP.

Beginning in privileged EXEC mode, follow these steps to enable loop guard. This procedure is optional.

	Command	Purpose
Step 1	show spanning-tree active or show spanning-tree mst	Determine which ports are alternate or root ports.
Step 2	configure terminal	Enter global configuration mode.
Step 3	spanning-tree loopguard default	Enable loop guard. By default, loop guard is disabled.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To globally disable loop guard, use the **no spanning-tree loopguard default** global configuration command. You can override the setting of the **no spanning-tree loopguard default** global configuration command by using the **spanning-tree guard loop** interface configuration command.

Displaying the Spanning-Tree Status

To display the spanning-tree status, use one or more of the privileged EXEC commands in [Table 16-2](#):

Table 16-2 *Commands for Displaying the Spanning-Tree Status*

Command	Purpose
show spanning-tree active	Displays spanning-tree information on active interfaces only.
show spanning-tree detail	Displays a detailed summary of interface information.
show spanning-tree interface <i>interface-id</i>	Displays spanning-tree information for the specified interface.
show spanning-tree mst interface <i>interface-id</i>	Displays MST information for the specified interface.
show spanning-tree summary [totals]	Displays a summary of port states or displays the total lines of the spanning-tree state section.

You can clear spanning-tree counters by using the **clear spanning-tree [interface *interface-id*]** privileged EXEC command.

For information about other keywords for the **show spanning-tree** privileged EXEC command, refer to the command reference for this release.



Configuring DHCP Features

This chapter describes how to configure Dynamic Host Configuration Protocol (DHCP) snooping and the option-82 data insertion features on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, see the command reference for this release, and refer to the “*IP Addressing and Services*” section in the *Cisco IOS IP and IP Routing Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding DHCP Features, page 17-1](#)
- [Configuring DHCP Features, page 17-3](#)
- [Displaying DHCP Information, page 17-5](#)

Understanding DHCP Features

DHCP is widely used in LAN environments to dynamically assign host IP addresses from a centralized server, which significantly reduces the overhead of administration of IP addresses. DHCP also helps conserve the limited IP address space because IP addresses no longer need to be permanently assigned to hosts; only those hosts that are connected to the network consume IP addresses.

DHCP Snooping

DHCP snooping is a DHCP security feature that provides network security by filtering untrusted DHCP messages and by building and maintaining a DHCP snooping binding table. An untrusted message is a message that is received from outside the network or firewall that can cause traffic attacks within your network.

The DHCP snooping binding table contains the MAC address, IP address, lease time, binding type, VLAN number, and interface information that corresponds to the local untrusted interfaces of a switch; it does not contain information regarding hosts interconnected with a trusted interface. An untrusted interface is an interface that is configured to receive messages from outside the network or firewall. A trusted interface is an interface that is configured to receive only messages from within the network.

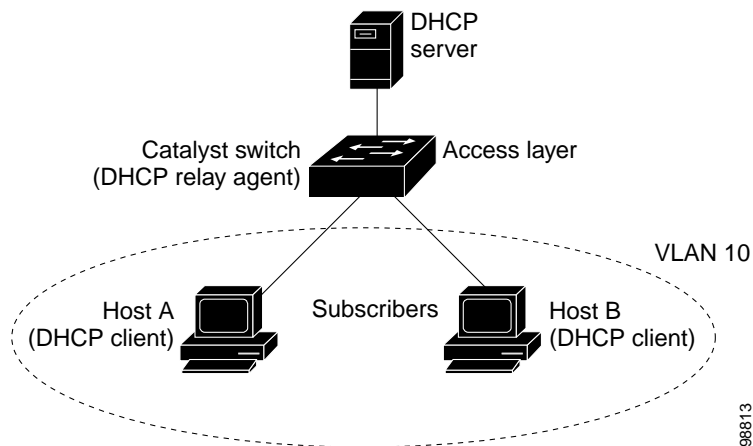
DHCP snooping acts like a firewall between untrusted hosts and DHCP servers. It also gives you a way to differentiate between untrusted interfaces connected to the end user and trusted interfaces connected to the DHCP server or another switch.

Option-82 Data Insertion

In residential, metropolitan Ethernet-access environments, DHCP can centrally manage the IP address assignments for a large number of subscribers. When the DHCP option-82 feature is enabled on the switch, a subscriber is identified by the switch port through which it connects to the network (in addition to its MAC address). Multiple hosts on the subscriber LAN can be connected to the same port on the access switch and are uniquely identified.

Figure 17-1 is an example of a metropolitan Ethernet network in which a centralized DHCP server assigns IP addresses to subscribers connected to the switch at the access layer. Because the DHCP clients and their associated DHCP server do not reside on the same IP network or subnet, a DHCP relay agent (the Catalyst switch) is configured with a helper address to enable broadcast forwarding and to transfer DHCP messages between the clients and the server.

Figure 17-1 DHCP Relay Agent in a Metropolitan Ethernet Network



When you enable the DHCP snooping information option 82 on the switch, this sequence of events occurs:

- The host (DHCP client) generates a DHCP request and broadcasts it on the network.
- When the switch receives the DHCP request, it adds the option-82 information in the packet. The option-82 information contains the switch MAC address (the remote ID suboption) and the port identifier, **vlan-mod-port**, from which the packet is received (the circuit ID suboption).
- The switch forwards the DHCP request that includes the option-82 field to the DHCP server.
- The DHCP server receives the packet. If the server is option-82-capable, it can use the remote ID, the circuit ID, or both to assign IP addresses and implement policies, such as restricting the number of IP addresses that can be assigned to a single remote ID or a circuit ID. Then the DHCP server echoes the option-82 field in the DHCP reply.
- The DHCP server unicasts the reply to the switch if the request was relayed to the server by the switch. The switch verifies that it originally inserted the option-82 data by inspecting the remote ID and possibly the circuit ID fields. The switch removes the option-82 field and forwards the packet to the switch port that connects to the DHCP client that sent the DHCP request.

Configuring DHCP Features

These sections describe how to configure DHCP snooping and option 82 on your switch:

- [Default DHCP Configuration, page 17-3](#)
- [DHCP Snooping Configuration Guidelines, page 17-3](#)
- [Enabling DHCP Snooping and Option 82, page 17-4](#)

Default DHCP Configuration

[Table 17-1](#) shows the default DHCP configuration.

Table 17-1 *Default DHCP Configuration*

Feature	Default Setting
DHCP snooping enabled globally	Disabled
DHCP snooping information option	Enabled
DHCP snooping limit rate	None configured
DHCP snooping trust	Untrusted
DHCP snooping VLAN	Disabled

DHCP Snooping Configuration Guidelines

These are the configuration guidelines for DHCP snooping.

- You must globally enable DHCP snooping on the switch.
- DHCP snooping is not active until DHCP snooping is enabled on a VLAN.
- When you globally enable DHCP snooping on the switch, these Cisco IOS commands are not available until snooping is disabled. If you enter these commands, the switch returns an error message, and the configuration is not applied.
 - **ip dhcp relay information check** global configuration command
 - **ip dhcp relay information policy** global configuration command
 - **ip dhcp relay information trust-all** global configuration command
 - **ip dhcp relay information trusted** interface configuration command
- Before configuring the DHCP information option on your switch, make sure to configure the device that is acting as the DHCP server. For example, you must specify the IP addresses that the DHCP server can assign or exclude or configure DHCP options for devices.
 - If your switch is the DHCP server, see the [“Configuring the DHCP Server” section on page 4-5](#) section for more information.
 - If your DHCP server is a Cisco device, refer to the “IP Addressing and Services” section in the “Configuring DHCP” chapter of the *Cisco IOS IP and IP Routing Configuration Guide for Release 12.1*. Otherwise, refer to the documentation that shipped with the server.

Enabling DHCP Snooping and Option 82

Beginning in privileged EXEC mode, follow these steps to enable DHCP snooping on the switch.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip dhcp snooping	Enable DHCP snooping globally.
Step 3	ip dhcp snooping vlan <i>vlan-id</i> [<i>vlan-id</i>]	Enable DHCP snooping on a VLAN or range of VLANs. You can specify a single VLAN identified by VLAN ID number or a start and end VLAN ID to specify a range of VLANs. The range is 1 to 4094.
Step 4	ip dhcp snooping information option	Enable the switch to insert and remove DHCP relay information (option-82 field) in forwarded DHCP request messages to the DHCP server. The default is enabled.
Step 5	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 6	ip dhcp snooping trust	(Optional) Configure the interface as trusted or untrusted. You can use the no keyword to configure an interface to receive messages from an untrusted client. The default is untrusted.
Step 7	ip dhcp snooping limit rate <i>rate</i>	(Optional) Configure the number of DHCP packets per second than an interface can receive. The range is 1 to 4294967294. The default is no rate limit configured. Note We recommend an untrusted rate limit of not more than 100 packets per second. Normally, the rate limit applies to untrusted interfaces. If you configure rate limiting for trusted interfaces, you will need to adjust the rate limit to a higher value because trusted interfaces might aggregate DHCP traffic in the switch.
Step 8	end	Return to privileged EXEC mode.
Step 9	show running-config	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable DHCP snooping, use the **no ip dhcp snooping** global configuration command. To disable DHCP snooping on a VLAN or range of VLANs, use the **no ip dhcp snooping vlan** *vlan-id* global configuration command. To disable the insertion and removal of the option-82 field, use the **no ip dhcp snooping information option** global configuration command.

This example shows how to enable DHCP snooping globally and on VLAN 10 and to configure a rate limit of 100 packets per second on Gigabit Ethernet port 0/1:

```
Switch(config)# ip dhcp snooping
Switch(config)# ip dhcp snooping vlan 10
Switch(config)# ip dhcp snooping information option
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip dhcp snooping limit rate 100
```

Displaying DHCP Information

You can display a DHCP snooping binding table and configuration information for all interfaces on a switch.

Displaying a Binding Table

The DHCP snooping binding table for each switch has binding entries that correspond to untrusted ports. The table does not have information about hosts interconnected with a trusted port because each interconnected switch has its own DHCP snooping binding table.

This example shows how to display the DHCP snooping binding entries for a switch.

```
Switch# show ip dhcp snooping binding
MacAddress      IpAddress      Lease(sec)  Type           VLAN  Interface
-----
00:30:94:C2:EF:35  41.0.0.51      286         dynamic        41    gigabitethernet0/3
00:D0:B7:1B:35:DE  41.0.0.52      237         dynamic        41    gigabitethernet0/3
00:00:00:00:00:01  40.0.0.46      286         dynamic        40    gigabitethernet0/4
00:00:00:00:00:03  42.0.0.33      286         dynamic        42    gigabitethernet0/4
00:00:00:00:00:02  41.0.0.53      286         dynamic        41    gigabitethernet0/4
```

Table 17-2 describes the fields in the `show ip dhcp snooping binding` command output.

Table 17-2 *show ip dhcp snooping binding Command Output*

Field	Description
MAC Address	Client hardware MAC address
IP Address	Client IP address assigned from the DHCP server
Lease (seconds)	IP address lease time
Type	Binding type; dynamic binding learned by DHCP snooping or statically configured binding
VLAN	VLAN number of the client interface
Interface	Interface that connects to the DHCP client host

Displaying the DHCP Snooping Configuration

This example shows how to display the DHCP snooping configuration for a switch.

```
Switch# show ip dhcp snooping
Switch DHCP snooping is enabled
DHCP snooping is configured on following VLANs:
40-42
Insertion of option 82 is enabled
Interface      Trusted      Rate limit (pps)
-----
gigabitethernet0/1    yes         unlimited
gigabitethernet0/2    yes         unlimited
gigabitethernet0/3    no          5000
gigabitethernet0/4    yes         unlimited
gigabitethernet0/7    yes         unlimited
gigabitethernet0/5    yes         unlimited
gigabitethernet0/7    yes         unlimited
```




Configuring IGMP Snooping and MVR

This chapter describes how to configure Internet Group Management Protocol (IGMP) snooping on the Catalyst 2970 switch, including an application of local IGMP snooping, Multicast VLAN Registration (MVR). It also includes procedures for controlling multicast group membership by using IGMP filtering and procedures for configuring the IGMP throttling action.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and the *Cisco IOS Release Network Protocols Command Reference, Part 1, for Release 12.1*.

This chapter consists of these sections:

- [Understanding IGMP Snooping, page 18-2](#)
- [Configuring IGMP Snooping, page 18-6](#)
- [Displaying IGMP Snooping Information, page 18-12](#)
- [Understanding Multicast VLAN Registration, page 18-13](#)
- [Configuring MVR, page 18-15](#)
- [Displaying MVR Information, page 18-19](#)
- [Configuring IGMP Filtering and Throttling, page 18-20](#)
- [Displaying IGMP Filtering and Throttling Configuration, page 18-25](#)



Note

You can either manage IP multicast group addresses through features such as IGMP snooping and MVR, or you can use static IP addresses.

Understanding IGMP Snooping

Layer 2 switches can use IGMP snooping to constrain the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. As the name implies, IGMP snooping requires the LAN switch to snoop on the IGMP transmissions between the host and the router and to keep track of multicast groups and member ports. When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry; when it receives an IGMP Leave Group message from a host, it removes the host port from the table entry. It also periodically deletes entries if it does not receive IGMP membership reports from the multicast clients.

**Note**

For more information on IP multicast and IGMP, refer to RFC 1112 and RFC 2236.

The multicast router sends out periodic general queries to all VLANs. All hosts interested in this multicast traffic send join requests and are added to the forwarding table entry. The switch creates one entry per VLAN in the IGMP snooping IP multicast forwarding table for each group from which it receives an IGMP join request.

The Catalyst 2970 switch supports IP multicast group-based bridging, rather than MAC-addressed based groups. With multicast MAC address-based groups, if an IP address being configured translates (aliases) to a previously configured MAC address or to any reserved multicast MAC addresses (in the range 224.0.0.xxx), the command fails. Because the Catalyst 2970 switch uses IP multicast groups, there are no address aliasing issues.

The IP multicast groups learned through IGMP snooping are dynamic. However, you can statically configure multicast groups by using the **ip igmp snooping vlan *vlan-id* static *ip_address* interface *interface-id*** global configuration command. If you specify group membership for a multicast group address statically, your setting supersedes any automatic manipulation by IGMP snooping. Multicast group membership lists can consist of both user-defined and IGMP snooping-learned settings.

If a port spanning-tree, a port group, or a VLAN ID change occurs, the IGMP snooping-learned multicast groups from this port on the VLAN are deleted.

These sections describe characteristics of IGMP snooping on the switch:

- [IGMP Versions, page 18-3](#)
- [Joining a Multicast Group, page 18-3](#)
- [Leaving a Multicast Group, page 18-5](#)
- [Immediate-Leave Processing, page 18-6](#)
- [IGMP Report Suppression, page 18-6](#)

IGMP Versions

The switch supports IGMP version 1, IGMP version 2, and IGMP version 3. These versions are interoperable on the switch. For example, if IGMP snooping is enabled on an IGMPv2 switch and the switch receives an IGMPv3 report from a host, the switch can forward the IGMPv3 report to the multicast router.

**Note**

The switches support IGMPv3 snooping based only on the destination multicast MAC address. They do not support snooping based on the source MAC address or on proxy reports.

An IGMPv3 switch supports Basic IGMPv3 Snooping Support (BISS), which includes support for the snooping features on IGMPv1 and IGMPv2 switches and for IGMPv3 membership report messages. BISS constrains the flooding of multicast traffic when your network includes IGMPv3 hosts. It constrains traffic to approximately the same set of ports as the IGMP snooping feature on IGMPv2 or IGMPv1 hosts.

**Note**

IGMPv3 join and leave messages are not supported on switches running IGMP filtering or MVR.

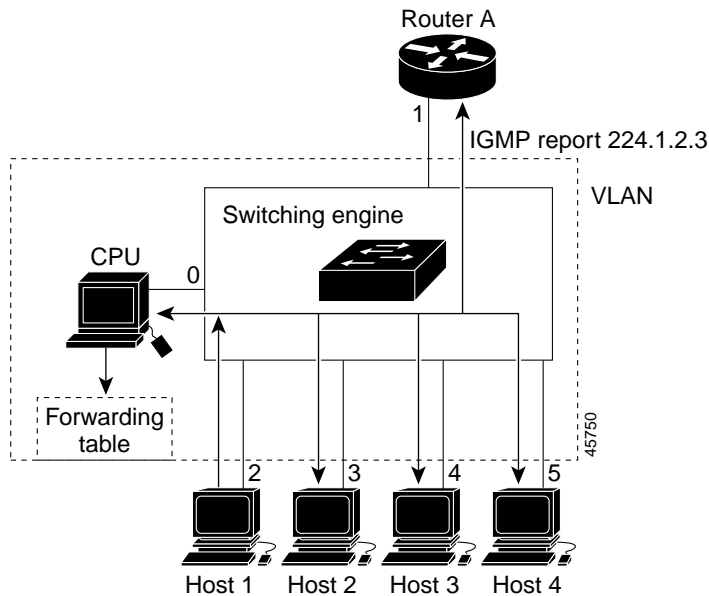
An IGMPv3 switch can receive messages from and forward messages to a device running the Source Specific Multicast (SSM) feature. For more information, refer to the “Configuring IP Multicast Layer 3 Switching” chapter in the *Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide, Cisco IOS Release 12.1(12c)EW* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat4000/12_1_12/config/mcastmls.htm

Joining a Multicast Group

When a host connected to the switch wants to join an IP multicast group, if it is an IGMP version 2 client, it sends an unsolicited IGMP join message, specifying the IP multicast group to join. Alternatively, when the switch receives a general query from the router, it forwards the query to all ports in the VLAN. IGMP version 1 or version 2 hosts wanting to join the multicast group respond by sending a join message to the switch. The switch CPU creates a multicast forwarding-table entry for the group if it is not already present. The CPU also adds the interface where the join message was received to the forwarding-table entry. The host associated with that interface receives multicast traffic for that multicast group. See [Figure 18-1](#).

Figure 18-1 Initial IGMP Join Message



Router A sends a general query to the switch, which forwards the query to ports 2 through 5, all members of the same VLAN. Host 1 wants to join multicast group 224.1.2.3 and multicasts an IGMP membership report (IGMP join message) to the group. When the CPU receives the IGMP report multicast by Host 1, the CPU uses the information in the IGMP report to set up a forwarding-table entry, as shown in [Table 18-1](#), that includes the port numbers connected to Host 1 and the router.

Table 18-1 IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
224.1.2.3	IGMP	1, 2

The switch hardware can distinguish IGMP information packets from other packets for the multicast group. The information in the table tells the switching engine to send frames addressed to the 224.1.2.3 multicast IP address that are not IGMP packets to the router and to the host that has joined the group.

If another host (for example, Host 4) sends an unsolicited IGMP join message for the same group ([Figure 18-2](#)), the CPU receives that message and adds the port number of Host 4 to the forwarding table as shown in [Table 18-2](#). Note that because the forwarding table directs IGMP messages to only the CPU, the message is not flooded to other ports on the switch. Any known multicast traffic is forwarded to the group and not to the CPU.

Figure 18-2 Second Host Joining a Multicast Group

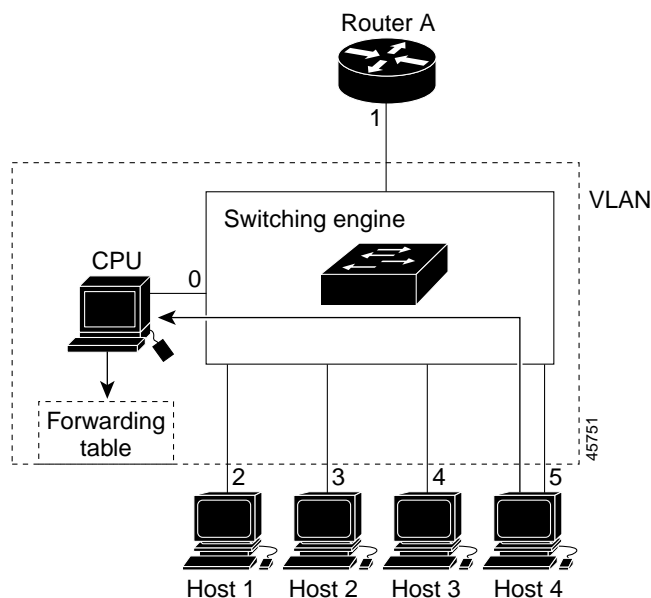


Table 18-2 Updated IGMP Snooping Forwarding Table

Destination Address	Type of Packet	Ports
224.1.2.3	IGMP	1, 2, 5

Leaving a Multicast Group

The router sends periodic multicast general queries, and the switch forwards these queries through all ports in the VLAN. Interested hosts respond to the queries. If at least one host in the VLAN wishes to receive multicast traffic, the router continues forwarding the multicast traffic to the VLAN. The switch forwards multicast group traffic only to those hosts listed in the forwarding table for that IP multicast group maintained by IGMP snooping.

When hosts want to leave a multicast group, they can silently leave or they can send a leave message. When the switch receives a leave message from a host, it sends out a MAC-based general query to determine if any other devices connected to that interface are interested in traffic for the specific multicast group. The switch then updates the forwarding table for that MAC group so that only those hosts interested in receiving multicast traffic for the group are listed in the forwarding table. If the router receives no reports from a VLAN, it removes the group for the VLAN from its IGMP cache.

Immediate-Leave Processing

Immediate Leave is only supported with IGMP version 2 hosts.

The switch uses IGMP snooping Immediate-Leave processing to remove from the forwarding table an interface that sends a leave message without the switch sending MAC-based general queries to the interface. The VLAN interface is pruned from the multicast tree for the multicast group specified in the original leave message. Immediate-Leave processing ensures optimal bandwidth management for all hosts on a switched network, even when multiple multicast groups are simultaneously in use.

**Note**

You should only use the Immediate-Leave processing feature on VLANs where a single host is connected to each port. If Immediate Leave is enabled in VLANs where more than one host is connected to a port, some hosts might inadvertently be dropped.

IGMP Report Suppression

**Note**

IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

The switch uses IGMP report suppression to forward only one IGMP report per multicast router query to multicast devices. When IGMP router suppression is enabled (the default), the switch sends the first IGMP report from all hosts for a group to all the multicast routers. The switch does not send the remaining IGMP reports for the group to the multicast routers. This feature prevents duplicate reports from being sent to the multicast devices.

If the multicast router query includes requests only for IGMPv1 and IGMPv2 reports, the switch forwards only the first IGMPv1 or IGMPv2 report from all hosts for a group to all the multicast routers.

If the multicast router query also includes requests for IGMPv3 reports, the switch forwards all IGMPv1, IGMPv2, and IGMPv3 reports for a group to the multicast devices.

If you disable IGMP report suppression, all IGMP reports are forwarded to the multicast routers.

Configuring IGMP Snooping

IGMP snooping allows switches to examine IGMP packets and make forwarding decisions based on their content.

These sections describe how to configure IGMP snooping:

- [Default IGMP Snooping Configuration, page 18-7](#)
- [Enabling or Disabling IGMP Snooping, page 18-7](#)
- [Setting the Snooping Method, page 18-8](#)
- [Configuring a Multicast Router Port, page 18-9](#)
- [Configuring a Host Statically to Join a Group, page 18-10](#)
- [Enabling IGMP Immediate-Leave Processing, page 18-10](#)
- [Disabling IGMP Report Suppression, page 18-11](#)

Default IGMP Snooping Configuration

Table 18-3 shows the default IGMP snooping configuration.

Table 18-3 Default IGMP Snooping Configuration

Feature	Default Setting
IGMP snooping	Enabled globally and per VLAN
Multicast routers	None configured
Multicast router learning (snooping) method	PIM-DVMRP
IGMP snooping Immediate Leave	Disabled
Static groups	None configured
IGMP report suppression	Enabled

Enabling or Disabling IGMP Snooping

By default, IGMP snooping is globally enabled on the switch. When globally enabled or disabled, it is also enabled or disabled in all existing VLAN interfaces. IGMP snooping is by default enabled on all VLANs, but can be enabled and disabled on a per-VLAN basis.

Global IGMP snooping overrides the VLAN IGMP snooping. If global snooping is disabled, you cannot enable VLAN snooping. If global snooping is enabled, you can enable or disable VLAN snooping.

Beginning in privileged EXEC mode, follow these steps to globally enable IGMP snooping on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping	Globally enable IGMP snooping in all existing VLAN interfaces.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To globally disable IGMP snooping on all VLAN interfaces, use the **no ip igmp snooping** global configuration command.

Beginning in privileged EXEC mode, follow these steps to enable IGMP snooping on a VLAN interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping vlan <i>vlan-id</i>	Enable IGMP snooping on the VLAN interface. The VLAN ID range is 1 to 4094. Note IGMP snooping must be globally enabled before you can enable VLAN snooping.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IGMP snooping on a VLAN interface, use the **no ip igmp snooping vlan *vlan-id*** global configuration command for the specified VLAN number.

Setting the Snooping Method

Multicast-capable router ports are added to the forwarding table for every Layer 2 multicast entry. The switch learns of such ports through one of these methods:

- Snooping on IGMP queries, Protocol Independent Multicast (PIM) packets, and Distance Vector Multicast Routing Protocol (DVMRP) packets
- Listening to Cisco Group Management Protocol (CGMP) packets from other routers
- Statically connecting to a multicast router port with the **ip igmp snooping mrouter** global configuration command

You can configure the switch either to snoop on IGMP queries and PIM/DVMRP packets or to listen to CGMP self-join or proxy-join packets. By default, the switch snoops on PIM/DVMRP packets on all VLANs. To learn of multicast router ports through only CGMP packets, use the **ip igmp snooping vlan *vlan-id* mrouter learn cgmp** global configuration command. When this command is entered, the router listens to only CGMP self-join and CGMP proxy-join packets and no other CGMP packets. To learn of multicast router ports through only PIM-DVMRP packets, use the **ip igmp snooping vlan *vlan-id* mrouter learn pim-dvmrp** global configuration command.



Note

If you want to use CGMP as the learning method and no multicast routers in the VLAN are CGMP proxy-enabled, you must enter the **ip cgmp router-only** command to dynamically access the router.

Beginning in privileged EXEC mode, follow these steps to alter the method in which a VLAN interface dynamically accesses a multicast router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping vlan <i>vlan-id</i> mrouter learn {cgmp pim-dvmrp}	Enable IGMP snooping on a VLAN. The VLAN ID range is 1 to 4094. Specify the multicast router learning method: <ul style="list-style-type: none"> • cgmp—Listen for CGMP packets. This method is useful for reducing control traffic. • pim-dvmrp—Snoop on IGMP queries and PIM-DVMRP packets. This is the default.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify the configuration.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to configure IGMP snooping to use CGMP packets as the learning method and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 mrouter learn cgmp
Switch(config)# end
Switch# show ip igmp snooping vlan 1
Global IGMP Snooping configuration:
-----
```



```

IGMP snooping           :Enabled
IGMPv3 snooping (minimal) :Enabled
Report suppression      :Enabled
TCN solicit query       :Disabled
TCN flood query count    :2

```

Vlan 1:

```

IGMP snooping           :Enabled
Immediate leave         :Disabled
Multicast router learning mode :pim-dvmrp
Source only learning age timer :10
CGMP interoperability mode :IGMP_ONLY

```

To return to the default learning method, use the **no ip igmp snooping vlan *vlan-id* mrouter learn cgmp** global configuration command.

Configuring a Multicast Router Port

To add a multicast router port (add a static connection to a multicast router), use the **ip igmp snooping vlan *vlan-id* mrouter interface *interface-id*** global configuration command on the switch.



Note

Static connections to multicast routers are supported only on switch ports.

Beginning in privileged EXEC mode, follow these steps to enable a static connection to a multicast router:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp snooping vlan <i>vlan-id</i> mrouter interface <i>interface-id</i>	Specify the multicast router VLAN ID and specify the interface to the multicast router. <ul style="list-style-type: none"> The VLAN ID range is 1 to 4094. The interface can be a physical interface or a port channel. The port channel range is 1 to 12.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping mrouter [vlan <i>vlan-id</i>]	Verify that IGMP snooping is enabled on the VLAN interface.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a multicast router port from the VLAN, use the **no ip igmp snooping vlan *vlan-id* mrouter interface *interface-id*** global configuration command.

This example shows how to enable a static connection to a multicast router and verify the configuration:

```

Switch# configure terminal
Switch(config)# ip igmp snooping vlan 200 mrouter interface gigabitethernet0/2
Switch(config)# end
Switch# show ip igmp snooping mrouter vlan 200
Vlan                ports
-----+-----
200                  Gi0/2(static)

```

Configuring a Host Statically to Join a Group

Hosts or Layer 2 ports normally join multicast groups dynamically, but you can also statically configure a host on an interface.

Beginning in privileged EXEC mode, follow these steps to add a Layer 2 port as a member of a multicast group:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip igmp snooping vlan <i>vlan-id</i> static <i>ip_address</i> interface <i>interface-id</i>	Statically configure a Layer 2 port as a member of a multicast group: <ul style="list-style-type: none"> • <i>vlan-id</i> is the multicast group VLAN ID. • <i>ip_address</i> is the group IP address. • <i>interface-id</i> is the member port. It can be a physical interface or port channel (1 to 12).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping multicast	Verify the member port and the IP address.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the Layer 2 port from the multicast group, use the **no ip igmp snooping vlan** *vlan-id* **static** *mac-address* **interface** *interface-id* global configuration command.

This example shows how to statically configure a host on an interface and verify the configuration:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 1 static 224.1.2.3 interface gigabitethernet0/1
Switch(config)# end
```

```
Switch# show ip igmp snooping multicast
```

Vlan	Group Address	Type	Ports
1	224.1.2.3	USER	Gi0/1

Enabling IGMP Immediate-Leave Processing

When you enable IGMP Immediate-Leave processing, the switch immediately removes a port when it detects an IGMP version 2 leave message on that port. You should use the Immediate-Leave feature only when there is a single receiver present on every port in the VLAN.



Note

Immediate Leave is supported with only IGMP version 2 hosts.

Beginning in privileged EXEC mode, follow these steps to enable IGMP Immediate-Leave processing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip igmp snooping vlan <i>vlan-id</i> immediate-leave	Enable IGMP Immediate-Leave processing on the VLAN interface.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping vlan <i>vlan-id</i>	Verify that Immediate Leave is enabled on the VLAN.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IGMP Immediate-Leave on a VLAN, use the **no ip igmp snooping vlan *vlan-id* immediate-leave** global configuration command.

This example shows how to enable IGMP Immediate-Leave processing on VLAN 130:

```
Switch# configure terminal
Switch(config)# ip igmp snooping vlan 130 immediate-leave
Switch(config)# end
```

Disabling IGMP Report Suppression



Note

IGMP report suppression is supported only when the multicast query has IGMPv1 and IGMPv2 reports. This feature is not supported when the query includes IGMPv3 reports.

IGMP report suppression is enabled by default. When it is enabled, the switch forwards only one IGMP report per multicast router query. When report suppression is disabled, all IGMP reports are forwarded to the multicast routers.

Beginning in privileged EXEC mode, follow these steps to disable IGMP report suppression:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no ip igmp snooping report-suppression	Disable IGMP report suppression.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip igmp snooping	Verify that IGMP report suppression is disabled.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To re-enable IGMP report suppression, use the **ip igmp snooping report-suppression** global configuration command.

Displaying IGMP Snooping Information

You can display IGMP snooping information for dynamically learned and statically configured router ports and VLAN interfaces. You can also display MAC address multicast entries for a VLAN configured for IGMP snooping.

To display IGMP snooping information, use one or more of the privileged EXEC commands in [Table 18-4](#).

Table 18-4 Commands for Displaying IGMP Snooping Information

Command	Purpose
show ip igmp snooping [vlan <i>vlan-id</i>]	Display the snooping configuration information for all VLANs on the switch or for a specified VLAN. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.
show ip igmp snooping multicast [count dynamic [count group <i>ip_address</i>] group <i>ip_address</i> user [count group <i>ip_address</i>]]	Display multicast table information for the switch or about a specific parameter: <ul style="list-style-type: none"> • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • group <i>ip_address</i>—Display characteristics of the multicast group with the specified group IP address. • user—Display only the user-configured multicast entries.
show ip igmp snooping multicast vlan <i>vlan-id</i> [count dynamic [count group <i>ip_address</i>] group <i>ip_address</i> user [count group <i>ip_address</i>]]	Display multicast table information for a multicast VLAN or about a specific parameter for the VLAN: <ul style="list-style-type: none"> • count—Display the total number of entries for the specified command options instead of the actual entries. • dynamic—Display entries learned through IGMP snooping. • group <i>ip_address</i>—Display characteristics of the multicast group with the specified group IP address. • user—Display only the user-configured multicast entries.
show ip igmp snooping mrouter [vlan <i>vlan-id</i>]	Display information on dynamically learned and manually configured multicast router interfaces. Note When you enable IGMP snooping, the switch automatically learns the interface to which a multicast router is connected. These are dynamically learned interfaces. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.
show ip igmp snooping querier [vlan <i>vlan-id</i>]	Display information about the IGMP version that an interface supports. (Optional) Enter vlan <i>vlan-id</i> to display information for a single VLAN.

For more information about the keywords and options in these commands, refer to the command reference for this release.

Understanding Multicast VLAN Registration

Multicast VLAN Registration (MVR) is designed for applications using wide-scale deployment of multicast traffic across an Ethernet ring-based service provider network (for example, the broadcast of multiple television channels over a service-provider network). MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on the network-wide multicast VLAN. It allows the single multicast VLAN to be shared in the network while subscribers remain in separate VLANs. MVR provides the ability to continuously send multicast streams in the multicast VLAN, but to isolate the streams from the subscriber VLANs for bandwidth and security reasons.

MVR assumes that subscriber ports subscribe and unsubscribe (join and leave) these multicast streams by sending out IGMP join and leave messages. These messages can originate from an IGMP version-2-compatible host with an Ethernet connection. Although MVR operates on the underlying mechanism of IGMP snooping, the two features operate independently of each other. One can be enabled or disabled without affecting the behavior of the other feature. However, if IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping.

The switch CPU identifies the MVR IP multicast streams and their associated IP multicast group in the switch forwarding table, intercepts the IGMP messages, and modifies the forwarding table to include or remove the subscriber as a receiver of the multicast stream, even though the receivers might be in a different VLAN from the source. This forwarding behavior selectively allows traffic to cross between different VLANs.

You can set the switch for compatible or dynamic mode of MVR operation.

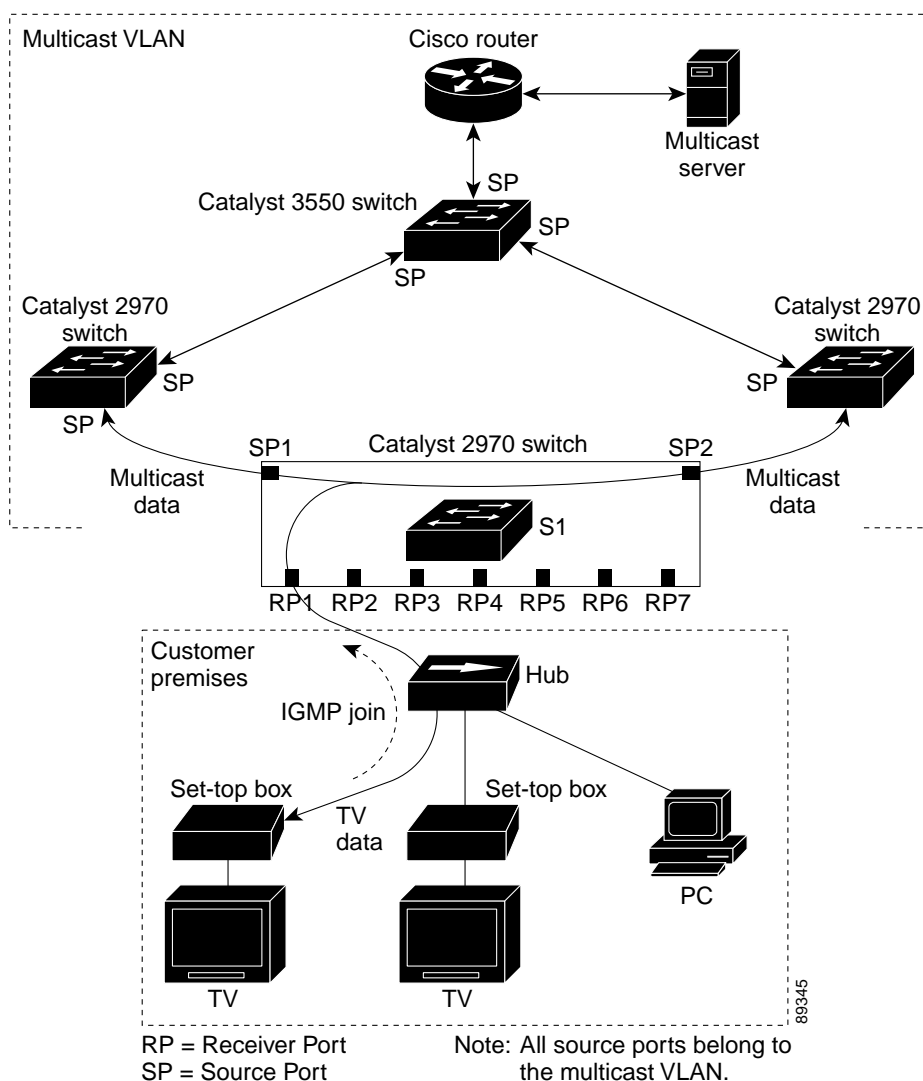
- In compatible mode, multicast data received by MVR hosts is forwarded to all MVR data ports, regardless of MVR host membership on those ports. The multicast data is forwarded only to those receiver ports which MVR hosts have explicitly joined, either by IGMP reports or by MVR static configuration. Also, IGMP reports received from MVR hosts are never forwarded out of MVR data ports that were configured in the switch.
- In dynamic mode, multicast data received by MVR hosts on the switch is forwarded from only those MVR data and client ports that the MVR hosts have explicitly joined, either by IGMP reports or by MVR static configuration. Any IGMP reports received from MVR hosts are also forwarded from all the MVR data ports in the switch. This eliminates using unnecessary bandwidth on MVR data port links, which occurs when the switch runs in compatible mode.

Only Layer 2 ports take part in MVR. You must configure ports as MVR receiver ports. Only one MVR multicast VLAN per switch is supported.

Using MVR in a Multicast Television Application

In a multicast television application, a PC or a television with a set-top box can receive the multicast stream. Multiple set-top boxes or PCs can be connected to one subscriber port, which is a switch port configured as an MVR receiver port. [Figure 18-3](#) is an example configuration. DHCP assigns an IP address to the set-top box or the PC. When a subscriber selects a channel, the set-top box or PC sends an IGMP report to the S1 switch to join the appropriate multicast. If the IGMP report matches one of the configured IP multicast group addresses, the switch CPU modifies the hardware address table to include this receiver port and VLAN as a forwarding destination of the specified multicast stream when it is received from the multicast VLAN. Uplink ports that send and receive multicast data to and from the multicast VLAN are called MVR source ports.

Figure 18-3 Multicast VLAN Registration Example



When a subscriber changes channels or turns off the television, the set-top box sends an IGMP leave message for the multicast stream. The switch CPU sends a MAC-based general query through the receiver port VLAN. If there is another set-top box in the VLAN still subscribing to this group, that set-top box must respond within the maximum response time specified in the query. If the CPU does not receive a response, it eliminates the receiver port as a forwarding destination for this group.

If the Immediate-Leave feature is enabled on a receiver port, the port leaves a multicast group more quickly. Without Immediate Leave, when the switch receives an IGMP leave message from a subscriber on a receiver port, it sends out an IGMP query on that port and waits for IGMP group membership reports. If no reports are received in a configured time period, the receiver port is removed from multicast group membership. With Immediate Leave, an IGMP query is not sent from the receiver port on which the IGMP leave was received. As soon as the leave message is received, the receiver port is removed from multicast group membership, which speeds up leave latency. Enable the Immediate Leave feature only on receiver ports to which a single receiver device is connected.

MVR eliminates the need to duplicate television-channel multicast traffic for subscribers in each VLAN. Multicast traffic for all channels is only sent around the VLAN trunk once—only on the multicast VLAN. The IGMP leave and join messages are in the VLAN to which the subscriber port is assigned. These messages dynamically register for streams of multicast traffic in the multicast VLAN on the Layer 3 device. The access layer switch (S1 switch) modifies the forwarding behavior to allow the traffic to be forwarded from the multicast VLAN to the subscriber port in a different VLAN, selectively allowing traffic to cross between two VLANs.

IGMP reports are sent to the same IP multicast group address as the multicast data. The S1 CPU must capture all IGMP join and leave messages from receiver ports and forward them to the multicast VLAN of the source (uplink) port, based on the MVR mode.

Configuring MVR

These sections include basic MVR configuration information:

- [Default MVR Configuration, page 18-16](#)
- [MVR Configuration Guidelines and Limitations, page 18-16](#)
- [Configuring MVR Global Parameters, page 18-16](#)
- [Configuring MVR Interfaces, page 18-18](#)

Default MVR Configuration

Table 18-5 shows the default MVR configuration.

Table 18-5 *Default MVR Configuration*

Feature	Default Setting
MVR	Disabled globally and per interface
Multicast addresses	None configured
Query response time	0.5 second
Multicast VLAN	VLAN 1
Mode	Compatible
Interface (per port) default	Neither a receiver nor a source port
Immediate Leave	Disabled on all ports

MVR Configuration Guidelines and Limitations

Follow these guidelines when configuring MVR:

- Receiver ports can only be access ports; they cannot be trunk ports. Receiver ports on a switch can be in different VLANs, but should not belong to the multicast VLAN.
- The maximum number of multicast entries (MVR group addresses) that can be configured on a switch (that is, the maximum number of television channels that can be received) is 256.
- MVR multicast data received in the source VLAN and leaving from receiver ports has its time-to-live (TTL) decremented by 1 in the Catalyst 2970 switch.
- Because MVR on the Catalyst 2970 switch uses IP multicast addresses instead of MAC multicast addresses, aliased IP multicast addresses are allowed on the switch. However, if the switch is interoperating with Catalyst 3550 or Catalyst 3500 XL switches, you should not configure IP addresses that alias between themselves or with the reserved IP multicast addresses (in the range 224.0.0.xxx).
- MVR can coexist with IGMP snooping on a switch.
- MVR data received on an MVR receiver port is not forwarded to MVR source ports.
- MVR does not support IGMPv3 messages.

Configuring MVR Global Parameters

You do not need to set the optional MVR parameters if you choose to use the default settings. If you do want to change the default parameters (except for the MVR VLAN), you must first enable MVR.



Note

For complete syntax and usage information for the commands used in this section, refer to the command reference for this release.

Beginning in privileged EXEC mode, follow these steps to configure MVR parameters:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mvr	Enable MVR on the switch.
Step 3	mvr group <i>ip-address</i> [<i>count</i>]	Configure an IP multicast address on the switch or use the <i>count</i> parameter to configure a contiguous series of MVR group addresses (the range for <i>count</i> is 1 to 256; the default is 1). Any multicast data sent to this address is sent to all source ports on the switch and all receiver ports that have elected to receive data on that multicast address. Each multicast address would correspond to one television channel.
Step 4	mvr querytime <i>value</i>	(Optional) Define the maximum time to wait for IGMP report memberships on a receiver port before removing the port from multicast group membership. The value is in units of tenths of a second. The range is from 1 to 100 and the default is 5 tenths or one-half second.
Step 5	mvr vlan <i>vlan-id</i>	(Optional) Specify the VLAN in which multicast data is received; all source ports must belong to this VLAN. The VLAN range is 1 to 4094. The default is VLAN 1.
Step 6	mvr mode { dynamic compatible }	(Optional) Specify the MVR mode of operation: <ul style="list-style-type: none"> • dynamic—Allows dynamic MVR membership on source ports. • compatible—Is compatible with Catalyst 3500 XL and Catalyst 2900 XL switches and does not support IGMP dynamic joins on source ports. The default is compatible mode.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mvr or show mvr members	Verify the configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the switch to its default settings, use the **no mvr** [**mode** | **group ip-address** | **querytime** | **vlan**] global configuration commands.

This example shows how to enable MVR, configure the group address, set the query time to 1 second (10 tenths), specify the MVR multicast VLAN as VLAN 22, and set the MVR mode as dynamic:

```
Switch(config)# mvr
Switch(config)# mvr group 228.1.23.4
Switch(config)# mvr querytime 10
Switch(config)# mvr vlan 22
Switch(config)# mvr mode dynamic
Switch(config)# end
```

You can use the **show mvr members** privileged EXEC command to verify the MVR multicast group addresses on the switch.

Configuring MVR Interfaces

Beginning in privileged EXEC mode, follow these steps to configure Layer 2 MVR interfaces:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mvr	Enable MVR on the switch.
Step 3	interface interface-id	Enter interface configuration mode, and enter the type and number of the Layer 2 port to configure.
Step 4	mvr type {source receiver}	<p>Configure an MVR port as one of these:</p> <ul style="list-style-type: none"> • source—Configure uplink ports that receive and send multicast data as source ports. Subscribers cannot be directly connected to source ports. All source ports on a switch belong to the single multicast VLAN. • receiver—Configure a port as a receiver port if it is a subscriber port and should only receive multicast data. It does not receive data unless it becomes a member of the multicast group, either statically or by using IGMP leave and join messages. Receiver ports cannot belong to the multicast VLAN. <p>The default configuration is as a non-MVR port. If you attempt to configure a non-MVR port with MVR characteristics, the operation fails.</p>
Step 5	mvr vlan vlan-id group [ip-address]	<p>(Optional) Statically configure a port to receive multicast traffic sent to the multicast VLAN and the IP multicast address. A port statically configured as a member of a group remains a member of the group until statically removed.</p> <p>Note In compatible mode, this command applies to only receiver ports. In dynamic mode, it applies to receiver ports and source ports.</p> <p>Receiver ports can also dynamically join multicast groups by using IGMP join and leave messages.</p>
Step 6	mvr immediate	<p>(Optional) Enable the Immediate Leave feature of MVR on the port.</p> <p>Note This command applies to only receiver ports and should only be enabled on receiver ports to which a single receiver device is connected.</p>
Step 7	end	Return to privileged EXEC mode.
Step 8	show mvr show mvr interface or show mvr members	Verify the configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to its default settings, use the **no mvr [type | immediate | vlan vlan-id | group]** interface configuration commands.

This example shows how to configure Gigabit Ethernet port 0/3 as a receiver port, statically configure the port to receive multicast traffic sent to the multicast group address, configure Immediate Leave on the interface, and verify the results.

```
Switch(config)# mvr
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# mvr type receiver
Switch(config-if)# mvr vlan 22 group 228.1.23.4
Switch(config-if)# mvr immediate
Switch(config)# end
Switch# show mvr interface
```

Port	Type	Status	Immediate Leave
Gi0/3	RECEIVER	ACTIVE/DOWN	ENABLED

Displaying MVR Information

You can display MVR information for the switch or for a specified interface. Beginning in privileged EXEC mode, use the commands in [Table 18-6](#) to display MVR configuration:

Table 18-6 Commands for Displaying MVR Information

Command	Purpose
show mvr	Displays MVR status and values for the switch—whether MVR is enabled or disabled, the multicast VLAN, the maximum (256) and current (0 through 256) number of multicast groups, the query response time, and the MVR mode.
show mvr interface <i>[interface-id]</i> [members <i>[vlan vlan-id]</i>]	<p>Displays all MVR interfaces and their MVR configurations.</p> <p>When a specific interface is entered, displays this information:</p> <ul style="list-style-type: none"> Type—Receiver or Source Status—One of these: <ul style="list-style-type: none"> Active means the port is part of a VLAN. Up/Down means that the port is forwarding or nonforwarding. Inactive means that the port is not part of any VLAN. Immediate Leave—Enabled or Disabled <p>If the members keyword is entered, displays all multicast group members on this port or, if a VLAN identification is entered, all multicast group members on the VLAN. The VLAN ID range is 1 to 4094; do not enter leading zeros.</p>
show mvr members <i>[ip-address]</i>	Displays all receiver and source ports that are members of any IP multicast group or the specified IP multicast group IP address.

Configuring IGMP Filtering and Throttling

In some environments, for example, metropolitan or multiple-dwelling unit (MDU) installations, you might want to control the set of multicast groups to which a user on a switch port can belong. You can control the distribution of multicast services, such as IP/TV, based on some type of subscription or service plan. You might also want to limit the number of multicast groups to which a user on a switch port can belong.

With the IGMP filtering feature, you can filter multicast joins on a per-port basis by configuring IP multicast profiles and associating them with individual switch ports. An IGMP profile can contain one or more multicast groups and specifies whether access to the group is permitted or denied. If an IGMP profile denying access to a multicast group is applied to a switch port, the IGMP join report requesting the stream of IP multicast traffic is dropped, and the port is not allowed to receive IP multicast traffic from that group. If the filtering action permits access to the multicast group, the IGMP report from the port is forwarded for normal processing.

IGMP filtering controls only group specific query and membership reports, including join and leave reports. It does not control general IGMP queries. IGMP filtering has no relationship with the function that directs the forwarding of IP multicast traffic. The filtering feature operates in the same manner whether CGMP or MVR is used to forward the multicast traffic.

IGMP filtering is only applicable to dynamic learning of IP multicast group addresses; not static configuration.

You can also set the maximum number of IGMP groups that a Layer 2 interface can join.

With the IGMP throttling feature, you can also set the maximum number of IGMP groups that a Layer 2 interface can join. If the maximum number of IGMP groups is set, the IGMP snooping forwarding table contains the maximum number of entries, and the interface receives an IGMP join report, you can configure an interface to drop the IGMP report or to remove a randomly selected multicast entry in the forwarding table and then to add the IGMP group in the report to the table.

**Note**

IGMPv3 join and leave messages are not supported on switches running IGMP filtering.

These sections describe how to configure IGMP filtering and throttling:

- [Default IGMP Filtering and Throttling Configuration, page 18-21](#)
- [Configuring IGMP Profiles, page 18-21](#) (optional)
- [Applying IGMP Profiles, page 18-22](#) (optional)
- [Setting the Maximum Number of IGMP Groups, page 18-23](#) (optional)
- [Configuring the IGMP Throttling Action, page 18-24](#) (optional)

Default IGMP Filtering and Throttling Configuration

Table 18-7 shows the default IGMP filtering configuration.

Table 18-7 Default IGMP Filtering Configuration

Feature	Default Setting
IGMP filters	None applied
IGMP maximum number of IGMP groups	No maximum set
IGMP profiles	None defined
IGMP profile action	Deny the range addresses

When the maximum number of groups is in forwarding table, the default IGMP throttling action is to deny the IGMP report. For configuration guidelines, see the [“Configuring the IGMP Throttling Action” section on page 18-24](#).

Configuring IGMP Profiles

To configure an IGMP profile, use the **ip igmp profile** global configuration command with a profile number to create an IGMP profile and to enter IGMP profile configuration mode. From this mode, you can specify the parameters of the IGMP profile to be used for filtering IGMP join requests from a port. When you are in IGMP profile configuration mode, you can create the profile by using these commands:

- **deny**: Specifies that matching addresses are denied; this is the default condition.
- **exit**: Exits from igmp-profile configuration mode.
- **no**: Negates a command or sets its defaults.
- **permit**: Specifies that matching addresses are permitted.
- **range**: Specifies a range of IP addresses for the profile. You can enter a single IP address or a range with a start and an end address.

The default is for the switch to have no IGMP profiles configured. When a profile is configured, if neither the **permit** nor **deny** keyword is included, the default is to deny access to the range of IP addresses.

Beginning in privileged EXEC mode, follow these steps to create an IGMP profile:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip igmp profile <i>profile number</i>	Enter IGMP profile configuration mode, and assign a number to the profile you are configuring. The range is from 1 to 4294967295.
Step 3	permit deny	(Optional) Set the action to permit or deny access to the IP multicast address. If no action is configured, the default for the profile is to deny access.

	Command	Purpose
Step 4	range <i>ip multicast address</i>	Enter the IP multicast address or range of IP multicast addresses to which access is being controlled. If entering a range, enter the low IP multicast address, a space, and the high IP multicast address. You can use the range command multiple times to enter multiple addresses or ranges of addresses.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip igmp profile <i>profile number</i>	Verify the profile configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete a profile, use the **no ip igmp profile** *profile number* global configuration command.

To delete an IP multicast address or range of IP multicast addresses, use the **no range ip multicast address** IGMP profile configuration command.

This example shows how to create IGMP profile 4 allowing access to the single IP multicast address and how to verify the configuration. If the action was to deny (the default), it would not appear in the **show ip igmp profile** output display.

```
Switch(config)# ip igmp profile 4
Switch(config-igmp-profile)# permit
Switch(config-igmp-profile)# range 229.9.9.0
Switch(config-igmp-profile)# end
Switch# show ip igmp profile 4
IGMP Profile 4
    permit
    range 229.9.9.0 229.9.9.0
```

Applying IGMP Profiles

To control access as defined in an IGMP profile, use the **ip igmp filter** interface configuration command to apply the profile to the appropriate interfaces. You can apply IGMP profiles to layer 2 access ports only. You cannot apply profiles to ports that belong to an EtherChannel port group. You can apply a profile to multiple interfaces, but each interface can only have one profile applied to it.

Beginning in privileged EXEC mode, follow these steps to apply an IGMP profile to a switch port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure. The interface must be a Layer 2 port that does not belong to an EtherChannel port group.
Step 3	ip igmp filter <i>profile number</i>	Apply the specified IGMP profile to the interface. The profile number can be from 1 to 4294967295.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a profile from an interface, use the **no ip igmp filter profile number** interface configuration command.

This example shows how to apply IGMP profile 4 to an interface:

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip igmp filter 4
Switch(config-if)# end
```

Setting the Maximum Number of IGMP Groups

You can set the maximum number of IGMP groups that a Layer 2 interface can join by using the **ip igmp max-groups** interface configuration command. Use the **no** form of this command to set the maximum back to the default, which is no limit.

You can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.

Beginning in privileged EXEC mode, follow these steps to set the maximum number of IGMP groups in the forwarding table:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the interface to configure. The interface can be a Layer 2 port that does not belong to an EtherChannel group or a EtherChannel interface.
Step 3	ip igmp max-groups <i>number</i>	Set the maximum number of IGMP groups that the interface can join. The range is from 0 to 4294967294. The default is to have no maximum set.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the maximum group limitation and return to the default of no maximum, use the **no ip igmp max-groups** interface configuration command.

This example shows how to limit to 25 the number of IGMP groups that an interface can join.

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip igmp max-groups 25
Switch(config-if)# end
```

Configuring the IGMP Throttling Action

After you set the maximum number of IGMP groups that a Layer 2 interface can join, you can configure an interface to remove a randomly selected multicast entry in the forwarding table and to add the next IGMP group to it by using the **ip igmp max-groups action replace** interface configuration command. Use the **no** form of this command to return to the default, which is to drop the IGMP join report.

Follow these guidelines when configuring the IGMP throttling action:

- You can use this command on a logical EtherChannel interface but cannot use it on ports that belong to an EtherChannel port group.
- When the maximum group limitation is set to the default (no maximum), entering the **ip igmp max-groups action {deny | replace}** command has no effect.
- If you configure the throttling action and set the maximum group limitation after an interface has added multicast entries to the forwarding table, the forwarding-table entries are either aged out or removed, depending on the throttling action.
 - If you configure the throttling action as **deny**, the entries that were previously in the forwarding table are not removed but are aged out. After these entries are aged out and the maximum number of entries is in the forwarding table, the switch drops the next IGMP report received on the interface.
 - If you configure the throttling action as **replace**, the entries that were previously in the forwarding table are removed. When the maximum number of entries is in the forwarding table, the switch deletes a randomly selected entry and adds an entry for the next IGMP report received on the interface.

To prevent the switch from removing the forwarding-table entries, you can configure the IGMP throttling action before an interface adds entries to the forwarding table.

Beginning in privileged EXEC mode, follow these steps to configure the throttling action when the maximum number of entries is in the forwarding table:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure. The interface can be a Layer 2 port that does not belong to an EtherChannel group or an EtherChannel interface. The interface cannot be a trunk port.
Step 3	ip igmp max-groups action {deny replace}	When an interface receives an IGMP report and the maximum number of entries is in the forwarding table, specify the action that the interface takes: <ul style="list-style-type: none"> • deny—Drop the report. • replace—Remove a randomly selected multicast entry in the forwarding table, and add the IGMP group in the report.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config interface <i>interface-id</i>	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default action of dropping the report, use the **no ip igmp max-groups action** interface configuration command.

This example shows how to configure an interface to remove a randomly selected multicast entry in the forwarding table and to add an IGMP group to the forwarding table when the maximum number of entries is in the table.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip igmp max-groups action replace
Switch(config-if)# end
```

Displaying IGMP Filtering and Throttling Configuration

You can display IGMP profile characteristics, and you can display the IGMP profile and maximum group configuration for all interfaces on the switch or for a specified interface. You can also display the IGMP throttling configuration for all interfaces on the switch or for a specified interface.

Use the privileged EXEC commands in [Table 18-8](#) to display IGMP filtering and throttling configuration:

Table 18-8 Commands for Displaying IGMP Filtering and Throttling Configuration

Command	Purpose
show ip igmp profile [<i>profile number</i>]	Displays the specified IGMP profile or all the IGMP profiles defined on the switch.
show running-config [<i>interface interface-id</i>]	Displays the configuration of the specified interface or the configuration of all interfaces on the switch, including (if configured) the maximum number of IGMP groups to which an interface can belong and the IGMP profile applied to the interface.



Configuring Port-Based Traffic Control

This chapter describes how to configure the port-based traffic control features on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Configuring Storm Control, page 19-1](#)
- [Configuring Protected Ports, page 19-5](#)
- [Configuring Port Blocking, page 19-6](#)
- [Configuring Port Security, page 19-7](#)
- [Displaying Port-Based Traffic Control Settings, page 19-15](#)

Configuring Storm Control

These sections include storm control configuration information and procedures:

- [Understanding Storm Control, page 19-2](#)
- [Default Storm Control Configuration, page 19-3](#)
- [Enabling Storm Control, page 19-3](#)

Understanding Storm Control

Storm control prevents switchports on a LAN from being disrupted by a broadcast, multicast, or unicast storm on one of the physical interfaces. A LAN storm occurs when packets flood the LAN, creating excessive traffic and degrading network performance. Errors in the protocol-stack implementation or in the network configuration can cause a storm.

Storm control (or traffic suppression) monitors incoming traffic statistics over a time period and compares the measurement with a predefined suppression level threshold. The threshold represents the percentage of the total available bandwidth of the port. The switch supports separate storm control thresholds for broadcast, multicast, and unicast traffic. If the threshold of a traffic type is reached, further traffic of that type is suppressed until the incoming traffic falls below the threshold level.



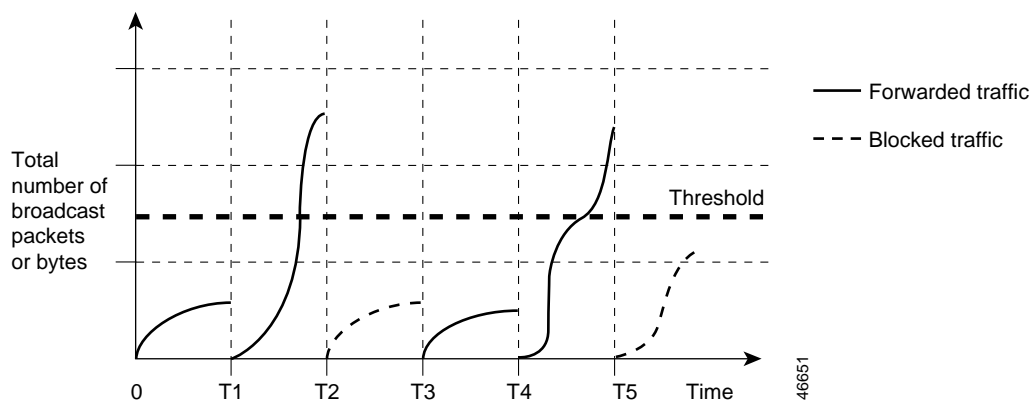
Note

When the storm control threshold for multicast traffic is reached, all multicast traffic except control traffic, such as bridge protocol data unit (BDPU) and Cisco Discovery Protocol (CDP) frames, are blocked.

When storm control is enabled, the switch monitors packets passing from an interface to the switching bus and determines if the packet is unicast, multicast, or broadcast. The switch monitors the number of broadcast, multicast, or unicast packets received within a 200-millisecond time interval, and when a threshold for one type of traffic is reached, that type of traffic is dropped. This threshold is specified as a percentage of total available bandwidth that can be used by broadcast (multicast or unicast) traffic.

The graph in [Figure 19-1](#) shows broadcast traffic patterns on an interface over a given period of time. The example can also be applied to multicast and unicast traffic. In this example, the broadcast traffic being forwarded exceeded the configured threshold between time intervals T1 and T2 and between T4 and T5. When the amount of specified traffic exceeds the threshold, all traffic of that kind is dropped for the next time period. Therefore, broadcast traffic is blocked during the intervals following T2 and T5. At the next time interval (for example, T3), if broadcast traffic does not exceed the threshold, it is again forwarded.

Figure 19-1 Broadcast Storm Control Example



The combination of the storm-control suppression level and the 200-millisecond time interval control the way the storm control algorithm works. A higher threshold allows more packets to pass through. A threshold value of 100 percent means that no limit is placed on the traffic. A value of 0.0 means that all broadcast, multicast, or unicast traffic on that port is blocked.

**Note**

Because packets do not arrive at uniform intervals, the 200-millisecond time interval during which traffic activity is measured can affect the behavior of storm control.

The switch continues to monitor traffic on the port, and when the utilization level is below the threshold level, the type of traffic that was dropped is forwarded again.

You use the **storm-control** interface configuration commands to set the threshold value for each traffic type.

**Note**

Although visible in the command-line interface (CLI) online help, the **switchport broadcast**, **switchport multicast**, and **switchport unicast** interface configuration commands for setting suppression levels are not available. These commands are obsolete, replaced by the **storm-control** interface configuration commands.

Default Storm Control Configuration

By default, unicast, broadcast, and multicast storm control is disabled on the switch interfaces; that is, the suppression level is 100 percent.

Enabling Storm Control

You enable storm control on an interface and enter the percentage of total available bandwidth that you want to be used by a particular type of traffic; entering 100 percent allows all traffic. However, because of hardware limitations and the way in which packets of different sizes are counted, threshold percentages are approximations. Depending on the sizes of the packets making up the incoming traffic, the actual enforced threshold might differ from the configured level by several percentage points.

**Note**

Storm control is supported only on physical interfaces; it is not supported on EtherChannel port channels even though the command is available in the CLI.

Beginning in privileged EXEC mode, follow these steps to enable a particular type of storm control:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the physical interface to configure, for example gigabitethernet0/1 .
Step 3	storm-control broadcast level <i>level</i> [<i>level</i>]	Specify the broadcast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all broadcast traffic on that port is blocked.

	Command	Purpose
Step 4	storm-control multicast level <i>level</i> [<i>.level</i>]	Specify the multicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all multicast traffic on that port is blocked.
Step 5	storm-control unicast level <i>level</i> [<i>.level</i>]	Specify the unicast traffic suppression level for an interface as a percentage of total bandwidth. The level can be from 1 to 100; the optional fraction of a level can be from 0 to 99. A threshold value of 100 percent means that no limit is placed on broadcast traffic. A value of 0.0 means that all unicast traffic on that port is blocked.
Step 6	end	Return to privileged EXEC mode.
Step 7	show storm-control [<i>interface-id</i>] [broadcast multicast unicast]	Verify the storm control suppression levels set on the interface for the specified traffic type. If you do not enter a traffic type, broadcast storm control settings are displayed.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable storm control, use the **no storm-control broadcast level**, **no storm-control multicast level**, or **no storm-control unicast level** interface configuration commands.

This example shows how to set the multicast storm control level at 70.5 percent on Gigabit Ethernet interface 0/1 and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# storm-control multicast level 70.5
Switch(config-if)# end
Switch# show storm-control gigabitethernet0/1 multicast
Interface  Filter State    Level    Current
-----
Gi0/1     Forwarding  70.50%   0.00%
```

This example shows how to disable the multicast storm control on Gigabit Ethernet interface 0/1 and verify the configuration:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# no storm-control multicast level
Switch(config-if)# end
Switch# show storm-control gigabitethernet0/1 multicast
Interface  Filter State    Level    Current
-----
Gi0/1     inactive   100.00%  N/A
```

Configuring Protected Ports

Some applications require that no traffic be forwarded at Layer 2 between ports on the same switch so that one neighbor does not see the traffic generated by another neighbor. In such an environment, the use of protected ports ensures that there is no exchange of unicast, broadcast, or multicast traffic between these ports on the switch.

Protected ports have these features:

- A protected port does not forward any traffic (unicast, multicast, or broadcast) to any other port that is also a protected port. Traffic cannot be forwarded between protected ports at Layer 2; all traffic passing between protected ports must be forwarded through a Layer 3 device.
- Forwarding behavior between a protected port and a nonprotected port proceeds as usual.

Default Protected Port Configuration

The default is to have no protected ports defined.

Protected Port Configuration Guidelines

You can configure protected ports on a physical interface (for example, Gigabit Ethernet 0/1) or an EtherChannel group (for example, port-channel 5). When you enable protected ports for a port channel, it is enabled for all ports in the port-channel group.

Configuring a Protected Port

Beginning in privileged EXEC mode, follow these steps to define a port as a protected port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the interface to configure, for example gigabitethernet0/1 .
Step 3	switchport protected	Configure the interface to be a protected port.
Step 4	end	Return to privileged EXEC mode.
Step 5	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable protected port, use the **no switchport protected** interface configuration command.

This example shows how to configure Gigabit Ethernet interface 0/1 as a protected port:

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport protected
Switch(config-if)# end
```

Configuring Port Blocking

By default, the switch floods packets with unknown destination MAC addresses out of all ports. If unknown unicast and multicast traffic is forwarded to a protected port, there could be security issues. To prevent unknown unicast or multicast traffic from being forwarded from one port to another, you can block a port (protected or nonprotected) from flooding unknown unicast or multicast packets to other ports.

Default Port Blocking Configuration

The default is to not block flooding of unknown multicast and unicast traffic out of a port, but to flood these packets to all ports.

Blocking Flooded Traffic on an Interface



Note

The interface can be a physical interface (for example, Gigabit Ethernet 0/1) or an EtherChannel group (for example, port-channel 5). When you block multicast or unicast traffic for a port channel, it is blocked on all ports in the port channel group.

Beginning in privileged EXEC mode, follow these steps to disable the flooding of multicast and unicast packets out of an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the type and number of the interface to configure, for example gigabitethernet0/1 .
Step 3	switchport block multicast	Block unknown multicast forwarding out of the port.
Step 4	switchport block unicast	Block unknown unicast forwarding out of the port.
Step 5	end	Return to privileged EXEC mode.
Step 6	show interfaces <i>interface-id</i> switchport	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to the default condition where no traffic is blocked and normal forwarding occurs on the port, use the **no switchport block {multicast | unicast}** interface configuration commands.

This example shows how to block unicast and multicast flooding on Gigabit Ethernet interface 0/1:

```
Switch# configure terminal
```



```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport block multicast
Switch(config-if)# switchport block unicast
Switch(config-if)# end
```

Configuring Port Security

You can use the port security feature to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port. When you assign secure MAC addresses to a secure port, the port does not forward packets with source addresses outside the group of defined addresses. If you limit the number of secure MAC addresses to one and assign a single secure MAC address, the workstation attached to that port is assured the full bandwidth of the port.

If a port is configured as a secure port and the maximum number of secure MAC addresses is reached, when the MAC address of a station attempting to access the port is different from any of the identified secure MAC addresses, a security violation occurs. Also, if a station with a secure MAC address configured or learned on one secure port attempts to access another secure port, a violation is flagged.

These sections include port security configuration information and procedures:

- [Understanding Port Security, page 19-7](#)
- [Default Port Security Configuration, page 19-9](#)
- [Configuration Guidelines, page 19-9](#)
- [Enabling and Configuring Port Security, page 19-10](#)
- [Enabling and Configuring Port Security Aging, page 19-13](#)

Understanding Port Security

This section contains information about these topics:

- [Secure MAC Addresses, page 19-7](#)
- [Security Violations, page 19-8](#)

Secure MAC Addresses

You configure the maximum number of secure addresses allowed on a port by using the **switchport port-security maximum** *value* interface configuration command.



Note

If you try to set the maximum value to a number less than the number of secure addresses already configured on an interface, the command is rejected.

The switch supports these types of secure MAC addresses:

- Static secure MAC addresses—These are manually configured by using the **switchport port-security mac-address** *mac-address* interface configuration command, stored in the address table, and added to the switch running configuration.
- Dynamic secure MAC addresses—These are dynamically configured, stored only in the address table, and removed when the switch restarts.

- *Sticky* secure MAC addresses—These can be dynamically learned or manually configured, stored in the address table, and added to the running configuration. If these addresses are saved in the configuration file, when the switch restarts, the interface does not need to dynamically reconfigure them.

You can configure an interface to convert the dynamic MAC addresses to sticky secure MAC addresses and to add them to the running configuration by enabling *sticky learning*. To enable sticky learning, enter the **switchport port-security mac-address sticky** interface configuration command. When you enter this command, the interface converts all the dynamic secure MAC addresses, including those that were dynamically learned before sticky learning was enabled, to sticky secure MAC addresses. All sticky secure MAC addresses are added to the running configuration.

The sticky secure MAC addresses do not automatically become part of the configuration file, which is the startup configuration used each time the switch restarts. If you save the sticky secure MAC addresses in the configuration file, when the switch restarts, the interface does not need to relearn these addresses. If you do not save the sticky secure addresses, they are lost.

If sticky learning is disabled, the sticky secure MAC addresses are converted to dynamic secure addresses and are removed from the running configuration.

The maximum number of secure MAC addresses that you can configure on a switch is determined by the maximum number of available MAC addresses allowed in the system. This number represents the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces.

Security Violations

It is a security violation when one of these situations occurs:

- The maximum number of secure MAC addresses have been added to the address table, and a station whose MAC address is not in the address table attempts to access the interface.
- An address learned or configured on one secure interface is seen on another secure interface in the same VLAN.

You can configure the interface for one of three violation modes, based on the action to be taken if a violation occurs:

- **protect**—when the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred.



Note

We do not recommend configuring the protect violation mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit.

- **restrict**—when the number of secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. In this mode, you are notified that a security violation has occurred. An SNMP trap is sent, a syslog message is logged, and the violation counter increments.
- **shutdown**—a port security violation causes the interface to become error-disabled and to shut down immediately, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. When a secure port is in the error-disabled state, you can bring it out

of this state by entering the **errdisable recovery cause psecure-violation** global configuration command, or you can manually re-enable it by entering the **shutdown** and **no shut down** interface configuration commands. This is the default mode.

Table 19-1 shows the violation mode and the actions taken when you configure an interface for port security.

Table 19-1 Security Violation Mode Actions

Violation Mode	Traffic is forwarded ¹	Sends SNMP trap	Sends syslog message	Displays error message ²	Violation counter increments	Shuts down port
protect	No	No	No	No	No	No
restrict	No	Yes	Yes	No	Yes	No
shutdown	No	Yes	Yes	No	Yes	Yes

1. Packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses.

2. The switch returns an error message if you manually configure an address that would cause a security violation.

Default Port Security Configuration

Table 19-2 shows the default port security configuration for an interface.

Table 19-2 Default Port Security Configuration

Feature	Default Setting
Port security	Disabled on a port.
Sticky address learning	Disabled.
Maximum number of secure MAC addresses per port	1.
Violation mode	Shutdown. The port shuts down when the maximum number of secure MAC addresses is exceeded.
Port security aging	Disabled. Aging time is 0. Static aging is disabled. Type is absolute.

Configuration Guidelines

Follow these guidelines when configuring port security:

- Port security can only be configured on static access ports or trunk ports. A secure port cannot be a dynamic access port.
- A secure port cannot be a destination port for Switched Port Analyzer (SPAN).
- A secure port cannot belong to a Gigabit EtherChannel port group.

- You cannot configure static secure or sticky secure MAC addresses in the voice VLAN.



Note Voice VLAN is only supported on access ports and not on trunk ports, even though the configuration is allowed.

- When you enable port security on an interface that is also configured with a voice VLAN, you must set the maximum allowed secure addresses on the port to two plus the maximum number of secure addresses allowed on the access VLAN. When the port is connected to a Cisco IP phone, the IP phone requires up to two MAC addresses. The IP phone address is learned on the voice VLAN and might also be learned on the access VLAN. Connecting a PC to the IP phone requires additional MAC addresses.
- If any type of port security is enabled on the access VLAN, dynamic port security is automatically enabled on the voice VLAN. You cannot configure port security on a per-VLAN basis.
- When a voice VLAN is configured on a secure port that is also configured as a sticky secure port, all addresses on the voice VLAN are learned as dynamic secure addresses, and all addresses seen on the access VLAN to which the port belongs are learned as sticky secure addresses.
- When you enter a maximum secure address value for an interface, and the new value is greater than the previous value, the new value overwrites the previously configured value. If the new value is less than the previous value and the number of configured secure addresses on the interface exceeds the new value, the command is rejected.
- The switch does not support port security aging of sticky secure MAC addresses.

Enabling and Configuring Port Security

Beginning in privileged EXEC mode, follow these steps to restrict input to an interface by limiting and identifying MAC addresses of the stations allowed to access the port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the physical interface to configure, for example gigabitethernet0/1 .
Step 3	switchport mode { access trunk }	Set the interface switchport mode as access or trunk; an interface in the default mode (dynamic auto) cannot be configured as a secure port.
Step 4	switchport port-security	Enable port security on the interface.

	Command	Purpose
Step 5	switchport port-security maximum <i>value</i> [vlan <i>vlan-list</i>]	<p>(Optional) Set the maximum number of secure MAC addresses for the interface. The maximum number of secure MAC addresses that you can configure on a switch is determined by the maximum number of available MAC addresses allowed in the system. This number represents the total of available MAC addresses, including those used for other Layer 2 functions and any other secure MAC addresses configured on interfaces.</p> <p>(Optional) For trunk ports, you can set the maximum number of secure MAC addresses on a VLAN. If the vlan keyword is not entered, the default value is used.</p> <ul style="list-style-type: none"> • vlan—set a per-VLAN maximum value. • vlan <i>vlan-list</i>—set a per-VLAN maximum value on a range of VLANs separated by a hyphen, or a series of VLANs separated by commas. For non-specified VLANs, the per-VLAN maximum value is used.
Step 6	switchport port-security violation { protect restrict shutdown }	<p>(Optional) Set the violation mode, the action to be taken when a security violation is detected, as one of these:</p> <ul style="list-style-type: none"> • protect—When the number of port secure MAC addresses reaches the maximum limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses to drop below the maximum value or increase the number of maximum allowable addresses. You are not notified that a security violation has occurred. <p>Note We do not recommend configuring the protect mode on a trunk port. The protect mode disables learning when any VLAN reaches its maximum limit, even if the port has not reached its maximum limit.</p> <ul style="list-style-type: none"> • restrict—When the number of secure MAC addresses reaches the limit allowed on the port, packets with unknown source addresses are dropped until you remove a sufficient number of secure MAC addresses or increase the number of maximum allowable addresses. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. • shutdown—The interface is error-disabled when a violation occurs, and the port LED turns off. An SNMP trap is sent, a syslog message is logged, and the violation counter increments. <p>Note When a secure port is in the error-disabled state, you can bring it out of this state by entering the errdisable recovery cause psecure-violation global configuration command, or you can manually re-enable it by entering the shutdown and no shutdown interface configuration commands.</p>
Step 7	switchport port-security mac-address <i>mac-address</i> [vlan <i>vlan-id</i>]	<p>(Optional) Enter a secure MAC address for the interface. You can use this command to enter the maximum number of secure MAC addresses. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned.</p> <p>(Optional) On a trunk port, you can specify the VLAN ID and the MAC address. If no VLAN ID is specified, the native VLAN is used.</p> <p>Note If you enable sticky learning after you enter this command, the secure addresses that were dynamically learned are converted to sticky secure MAC addresses and are added to the running configuration.</p>

	Command	Purpose
Step 8	switchport port-security mac-address sticky	(Optional) Enable stick learning on the interface.
Step 9	switchport port-security mac-address sticky <i>mac-address</i>	(Optional) Enter a sticky secure MAC address, repeating the command as many times as necessary. If you configure fewer secure MAC addresses than the maximum, the remaining MAC addresses are dynamically learned, are converted to sticky secure MAC addresses, and are added to the running configuration. Note If you do not enable sticky learning before this command is entered, an error message appears, and you cannot enter a sticky secure MAC address.
Step 10	end	Return to privileged EXEC mode.
Step 11	show port-security	Verify your entries.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the interface to the default condition as not a secure port, use the **no switchport port-security** interface configuration command. If you enter this command when sticky learning is enabled, the sticky secure addresses remain part of the running configuration but are removed from the address table. All addresses are now dynamically learned.

To return the interface to the default number of secure MAC addresses, use the **no switchport port-security maximum value** interface configuration command. To return the violation mode to the default condition (shutdown mode), use the **no switchport port-security violation {protocol | restrict}** interface configuration command.

To disable sticky learning on an interface, use the **no switchport port-security mac-address sticky** interface configuration command. The interface converts the sticky secure MAC addresses to dynamic secure addresses. However, if you have previously saved the configuration with the sticky MAC addresses, you should save the configuration again after entering the **no switchport port-security mac-address sticky** command, or the sticky addresses will be restored if the switch reboots.

To delete a specific secure MAC address from the address table, use the **no switchport port-security mac-address mac-address** interface configuration command.

To delete all dynamic secure addresses on an interface from the address table, enter the **no switchport port-security** interface configuration command followed by the **switchport port-security** command (to re-enable port security on the interface). If you use the **no switchport port-security mac-address sticky** interface configuration command to convert sticky secure MAC addresses to dynamic secure MAC addresses before entering the **no switchport port-security** command, all secure addresses on the interface except those that were manually configured are deleted.

You must specifically delete configured secure MAC addresses from the address table by using the **no switchport port-security mac-address mac-address** interface configuration command.

This example shows how to enable port security on Gigabit Ethernet port 0/1 and to set the maximum number of secure addresses to 50. The violation mode is the default, no static secure MAC addresses are configured, and sticky learning is enabled.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport mode access
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security maximum 50
Switch(config-if)# switchport port-security mac-address sticky
```

This example shows how to configure a static secure MAC address on VLAN 3 on Gigabit Ethernet port 0/2:

```
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport port-security
Switch(config-if)# switchport port-security mac-address 0000.02000.0004 vlan 3
```

Enabling and Configuring Port Security Aging

You can use port security aging to set the aging time for all secure addresses on a port. Two types of aging are supported per port:

- Absolute—The secure addresses on the port are deleted after the specified aging time.
- Inactivity—The secure addresses on the port are deleted only if the secure addresses are inactive for the specified aging time.

Use this feature to remove and add devices on a secure port without manually deleting the existing secure MAC addresses and to still limit the number of secure addresses on a port. You can enable or disable the aging of secure addresses on a per-port basis.

Beginning in privileged EXEC mode, follow these steps to configure port security aging:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode for the port on which you want to enable port security aging.

	Command	Purpose
Step 3	switchport port-security aging {static time <i>time</i> type {absolute inactivity}}	<p>Enable or disable static aging for the secure port, or set the aging time or type.</p> <p>Note The switch does not support port security aging of sticky secure addresses.</p> <p>Enter static to enable aging for statically configured secure addresses on this port.</p> <p>For <i>time</i>, specify the aging time for this port. The valid range is from 0 to 1440 minutes. If the time is equal to 0, aging is disabled for this port.</p> <p>For type, select one of these keywords:</p> <ul style="list-style-type: none"> • absolute—Sets the aging type as absolute aging. All the secure addresses on this port age out exactly after the time (minutes) specified lapses and are removed from the secure address list. • inactivity—Sets the aging type as inactivity aging. The secure addresses on this port age out only if there is no data traffic from the secure source addresses for the specified time period.
Step 4	end	Return to privileged EXEC mode.
Step 5	show port-security [interface <i>interface-id</i>] [address]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable port security aging for all secure addresses on a port, use the **no switchport port-security aging time** interface configuration command. To disable aging for only statically configured secure addresses, use the **no switchport port-security aging static** interface configuration command.

This example shows how to set the aging time as 2 hours for the secure addresses on Gigabit Ethernet interface 0/1:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# switchport port-security aging time 120
```

This example shows how to set the aging time as 2 minutes for the inactivity aging type with aging enabled for the configured secure addresses on the interface:

```
Switch(config-if)# switchport port-security aging time 2
Switch(config-if)# switchport port-security aging type inactivity
Switch(config-if)# switchport port-security aging static
```

You can verify the previous commands by entering the **show port-security interface *interface-id*** privileged EXEC command.

Displaying Port-Based Traffic Control Settings

The **show interfaces *interface-id* switchport** privileged EXEC command displays (among other characteristics) the interface traffic suppression and control configuration. The **show interfaces counters** privileged EXEC commands display the count of discarded packets. The **show storm-control** and **show port-security** privileged EXEC commands display those features.

To display traffic control information, use one or more of the privileged EXEC commands in [Table 19-3](#).

Table 19-3 Commands for Displaying Traffic Control Status and Configuration

Command	Purpose
show interfaces [<i>interface-id</i>] switchport	Displays the administrative and operational status of all switching (nonrouting) ports or the specified port, including port blocking and port protection settings.
show storm-control [<i>interface-id</i>] [broadcast multicast unicast]	Displays storm control suppression levels set on all interfaces or the specified interface for the specified traffic type or for broadcast traffic if no traffic type is entered.
show interfaces [<i>interface-id</i>] counters broadcast	Displays the storm-control broadcast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [<i>interface-id</i>] counters multicast	Displays the storm-control multicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show interfaces [<i>interface-id</i>] counters unicast	Displays the storm-control unicast suppression discard counter with the number of packets discarded for all interfaces or the specified interface.
show port-security [interface <i>interface-id</i>]	Displays port security settings for the switch or for the specified interface, including the maximum allowed number of secure MAC addresses for each interface, the number of secure MAC addresses on the interface, the number of security violations that have occurred, and the violation mode.
show port-security [interface <i>interface-id</i>] address	Displays all secure MAC addresses configured on all switch interfaces or on a specified interface with aging information for each address.
show port-security interface <i>interface-id</i> vlan	Displays the number of secure MAC addresses configured per VLAN on the specified interface.



Configuring CDP

This chapter describes how to configure Cisco Discovery Protocol (CDP) on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding CDP, page 20-1](#)
- [Configuring CDP, page 20-2](#)
- [Monitoring and Maintaining CDP, page 20-5](#)

Understanding CDP

CDP is a device discovery protocol that runs over Layer 2 (the data link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches) and allows network management applications to discover Cisco devices that are neighbors of already known devices. With CDP, network management applications can learn the device type and the Simple Network Management Protocol (SNMP) agent address of neighboring devices running lower-layer, transparent protocols. This feature enables applications to send SNMP queries to neighboring devices.

CDP runs on all media that support Subnetwork Access Protocol (SNAP). Because CDP runs over the data-link layer only, two systems that support different network-layer protocols can learn about each other.

Each CDP-configured device sends periodic messages to a multicast address, advertising at least one address at which it can receive SNMP messages. The advertisements also contain time-to-live, or holdtime information, which is the length of time a receiving device holds CDP information before discarding it. Each device also listens to the messages sent by other devices to learn about neighboring devices.

On the switch, CDP enables the Cluster Management Suite to display a graphical view of the network. The switch uses CDP to find cluster candidates and maintain information about cluster members and other devices up to three cluster-enabled devices away from the command switch by default.

The switch supports CDP version 2.

Configuring CDP

These sections include CDP configuration information and procedures:

- [Default CDP Configuration, page 20-2](#)
- [Configuring the CDP Characteristics, page 20-2](#)
- [Disabling and Enabling CDP, page 20-3](#)
- [Disabling and Enabling CDP on an Interface, page 20-4](#)

Default CDP Configuration

Table 20-1 shows the default CDP configuration.

Table 20-1 Default CDP Configuration

Feature	Default Setting
CDP global state	Enabled
CDP interface state	Enabled
CDP timer (packet update frequency)	60 seconds
CDP holdtime (before discarding)	180 seconds
CDP version-2 advertisements	Enabled

Configuring the CDP Characteristics

You can configure the frequency of CDP updates, the amount of time to hold the information before discarding it, and whether or not to send version-2 advertisements.

Beginning in privileged EXEC mode, follow these steps to configure the CDP timer, holdtime, and advertisement type.

**Note**

Steps 2 through 4 are all optional and can be performed in any order.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp timer <i>seconds</i>	(Optional) Set the transmission frequency of CDP updates in seconds. The range is from 5 to 254; the default is 60 seconds.
Step 3	cdp holdtime <i>seconds</i>	(Optional) Specify the amount of time a receiving device should hold the information sent by your device before discarding it. The range is from 10 to 255 seconds; the default is 180 seconds.
Step 4	cdp advertise-v2	(Optional) Configure CDP to send version-2 advertisements. This is the default state.
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show cdp	Verify configuration by displaying global information about CDP on the device.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of the CDP commands to return to the default settings.

This example shows how to configure and verify CDP characteristics.

```
Switch# configure terminal
Switch(config)# cdp timer 50
Switch(config)# cdp holdtime 120
Switch(config)# cdp advertise-v2
Switch(config)# end

Switch# show cdp
Global CDP information:
  Sending CDP packets every 50 seconds
  Sending a holdtime value of 120 seconds
  Sending CDPv2 advertisements is enabled
```

For additional CDP **show** commands, see the [“Monitoring and Maintaining CDP”](#) section on page 20-5.

Disabling and Enabling CDP

CDP is enabled by default.



Note

Switch clusters and other Cisco devices (such as Cisco IP Phones) regularly exchange CDP messages. Disabling CDP can interrupt cluster discovery and device connectivity. For more information, see [Chapter 5, “Clustering Switches.”](#)

Beginning in privileged EXEC mode, follow these steps to disable the CDP device discovery capability:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no cdp run	Disable CDP.
Step 3	end	Return to privileged EXEC mode.

Beginning in privileged EXEC mode, follow these steps to enable CDP when it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp run	Enable CDP after disabling it.
Step 3	end	Return to privileged EXEC mode.

This example shows how to enable CDP if it has been disabled.

```
Switch# configure terminal
Switch(config)# cdp run
Switch(config)# end
```

Disabling and Enabling CDP on an Interface

CDP is enabled by default on all supported interfaces to send and receive CDP information.

Beginning in privileged EXEC mode, follow these steps to disable CDP on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Enter interface configuration mode, and enter the interface on which you are disabling CDP.
Step 3	no cdp enable	Disable CDP on an interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Beginning in privileged EXEC mode, follow these steps to enable CDP on an interface when it has been disabled:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Enter interface configuration mode, and enter the interface on which you are enabling CDP.
Step 3	cdp enable	Enable CDP on an interface after disabling it.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

This example shows how to enable CDP on an interface when it has been disabled.

```
Switch# configure terminal
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# cdp enable
Switch(config-if)# end
```

Monitoring and Maintaining CDP

To monitor and maintain CDP on your device, perform one or more of these tasks, beginning in privileged EXEC mode.

Command	Description
clear cdp counters	Reset the traffic counters to zero.
clear cdp table	Delete the CDP table of information about neighbors.
show cdp	Display global information, such as frequency of transmissions and the holdtime for packets being sent.
show cdp entry <i>entry-name</i> [protocol version]	Display information about a specific neighbor. You can enter an asterisk (*) to display all CDP neighbors, or you can enter the name of the neighbor about which you want information. You can also limit the display to information about the protocols enabled on the specified neighbor or information about the version of software running on the device.
show cdp interface [<i>type number</i>]	Display information about interfaces where CDP is enabled. You can limit the display to the type of interface or the number of the interface about which you want information (for example, entering gigabitethernet 0/1 displays information only about Gigabit Ethernet port 1).
show cdp neighbors [<i>type number</i>] [detail]	Display information about neighbors, including device type, interface type and number, holdtime settings, capabilities, platform, and port ID. You can limit the display to neighbors on a specific type or number of interface or expand the display to provide more detailed information.
show cdp traffic	Display CDP counters, including the number of packets sent and received and checksum errors.

This is an example of the output from the **show cdp** privileged EXEC commands:

```
Switch# show cdp
Global CDP information:
  Sending CDP packets every 50 seconds
  Sending a holdtime value of 120 seconds
  Sending CDPv2 advertisements is enabled
```




Configuring UDLD

This chapter describes how to configure the UniDirectional Link Detection (UDLD) protocol on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding UDLD, page 21-1](#)
- [Configuring UDLD, page 21-4](#)
- [Displaying UDLD Status, page 21-6](#)

Understanding UDLD

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it administratively shuts down the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

Modes of Operation

UDLD supports two modes of operation: normal (the default) and aggressive. In normal mode, UDLD can detect unidirectional links due to misconnected interfaces on fiber-optic connections. In aggressive mode, UDLD can also detect unidirectional links due to one-way traffic on fiber-optic and twisted-pair links and to misconnected interfaces on fiber-optic links.

In normal and aggressive modes, UDLD works with the Layer 1 mechanisms to determine the physical status of a link. At Layer 1, autonegotiation takes care of physical signaling and fault detection. UDLD performs tasks that autonegotiation cannot perform, such as detecting the identities of neighbors and shutting down misconnected interfaces. When you enable both autonegotiation and UDLD, the Layer 1 and Layer 2 detections work together to prevent physical and logical unidirectional connections and the malfunctioning of other protocols.

A unidirectional link occurs whenever traffic sent by a local device is received by its neighbor but traffic from the neighbor is not received by the local device.

In normal mode, UDLD detects a unidirectional link when fiber strands in a fiber-optic interface are misconnected and the Layer 1 mechanisms do not detect this misconnection. If the interfaces are connected correctly but the traffic is one way, UDLD does not detect the unidirectional link because the Layer 1 mechanism, which is supposed to detect this condition, does not do so. In case, the logical link is considered undetermined, and UDLD does not disable the interface.

When UDLD is in normal mode, if one of the fiber strands in a pair is disconnected and autonegotiation is active, the link does not stay up because the Layer 1 mechanisms did not detect a physical problem with the link. In this case, UDLD does not take any action, and the logical link is considered undetermined.

In aggressive mode, UDLD detects a unidirectional link by using the previous detection methods. UDLD in aggressive mode can also detect a unidirectional link on a point-to-point link on which no failure between the two devices is allowed. It can also detect a unidirectional link when one of these problems exists:

- On fiber-optic or twisted-pair links, one of the interfaces cannot send or receive traffic.
- On fiber-optic or twisted-pair links, one of the interfaces is down while the other is up.
- One of the fiber strands in the cable is disconnected.

In these cases, UDLD shuts down the affected interface.

In a point-to-point link, UDLD hello packets can be considered as a heart beat whose presence guarantees the health of the link. Conversely, the loss of the heart beat means that the link must be shut down if it is not possible to re-establish a bidirectional link.

If both fiber strands in a cable are working normally from a Layer 1 perspective, UDLD in aggressive mode determines whether those fiber strands are connected correctly and whether traffic is flowing bidirectionally between the correct neighbors. This check cannot be performed by autonegotiation because autonegotiation operates at Layer 1.

Methods to Detect Unidirectional Links

UDLD operates by using two mechanisms:

- Neighbor database maintenance

UDLD learns about other UDLD-capable neighbors by periodically sending a hello packet (also called an advertisement or probe) on every active interface to keep each device informed about its neighbors.

When the switch receives a hello message, it caches the information until the age time (hold time or time-to-live) expires. If the switch receives a new hello message before an older cache entry ages, the switch replaces the older entry with the new one.

Whenever an interface is disabled and UDLD is running, whenever UDLD is disabled on an interface, or whenever the switch is reset, UDLD clears all existing cache entries for the interfaces affected by the configuration change. UDLD sends at least one message to inform the neighbors to flush the part of their caches affected by the status change. The message is intended to keep the caches synchronized.

- Event-driven detection and echoing

UDLD relies on echoing as its detection mechanism. Whenever a UDLD device learns about a new neighbor or receives a resynchronization request from an out-of-sync neighbor, it restarts the detection window on its side of the connection and sends echo messages in reply. Because this behavior is the same on all UDLD neighbors, the sender of the echoes expects to receive an echo in reply.

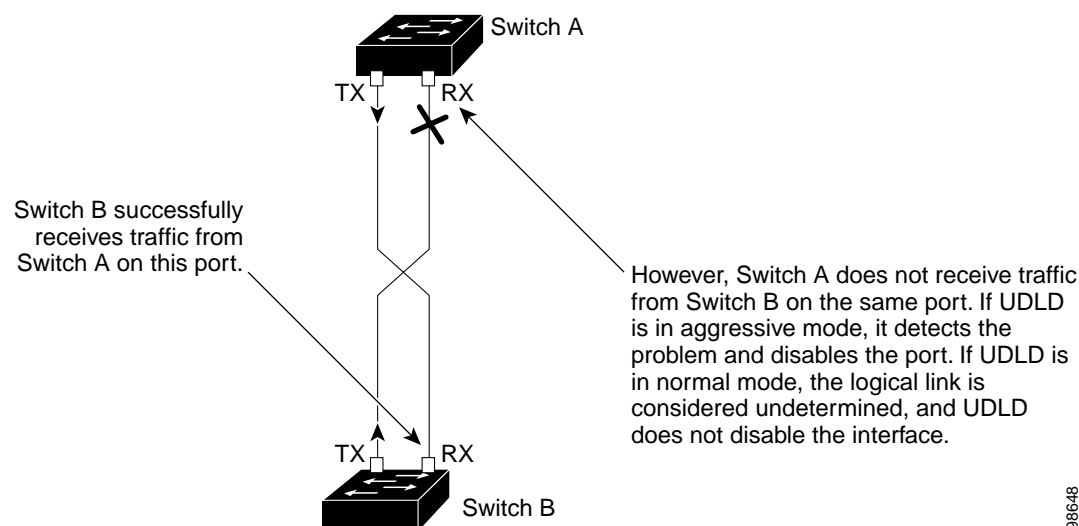
If the detection window ends and no valid reply message is received, the link might shut down, depending on the UDLD mode. When UDLD is in normal mode, the link might be considered undetermined and might not be shut down. When UDLD is in aggressive mode, the link is considered unidirectional, and the interface is shut down.

If UDLD in normal mode is in the advertisement or in the detection phase and all the neighbor cache entries are aged out, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbors.

If you enable aggressive mode when all the neighbors of a port have aged out either in the advertisement or in the detection phase, UDLD restarts the link-up sequence to resynchronize with any potentially out-of-sync neighbor. UDLD shuts down the port if, after the fast train of messages, the link state is still undetermined.

Figure 21-1 shows an example of a unidirectional link condition.

Figure 21-1 UDLD Detection of a Unidirectional Link



98648

Configuring UDLD

This section describes how to configure UDLD on your switch. It contains this configuration information:

- [Default UDLD Configuration, page 21-4](#)
- [Configuration Guidelines, page 21-4](#)
- [Enabling UDLD Globally, page 21-5](#)
- [Enabling UDLD on an Interface, page 21-5](#)
- [Resetting an Interface Disabled by UDLD, page 21-6](#)

Default UDLD Configuration

[Table 21-1](#) shows the default UDLD configuration.

Table 21-1 Default UDLD Configuration

Feature	Default Setting
UDLD global enable state	Globally disabled
UDLD per-interface enable state for fiber-optic media	Disabled on all Ethernet fiber-optic interfaces
UDLD per-interface enable state for twisted-pair (copper) media	Disabled on all Ethernet 10/100 and 1000BASE-TX interfaces
UDLD aggressive mode	Disabled

Configuration Guidelines

These are the UDLD configuration guidelines:

- UDLD is not supported on ATM interfaces.
- A UDLD-capable interface also cannot detect a unidirectional link if it is connected to a UDLD-incapable port of another switch.
- When configuring the mode (normal or aggressive), make sure that the same mode is configured on both sides of the link.

Enabling UDLD Globally

Beginning in privileged EXEC mode, follow these steps to enable UDLD in the aggressive or normal mode and to set the configurable message timer on all fiber-optic interfaces on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	udld { aggressive enable message time message-timer-interval }	<p>Specify the UDLD mode of operation:</p> <ul style="list-style-type: none"> • aggressive—Enables UDLD in aggressive mode on all fiber-optic interfaces. • enable—Enables UDLD in normal mode on all fiber-optic interfaces on the switch. UDLD is disabled by default. An individual interface configuration overrides the setting of the udld enable global configuration command. For more information about aggressive and normal modes, see the “Modes of Operation” section on page 21-1. • message time message-timer-interval—Configures the period of time between UDLD probe messages on ports that are in the advertisement phase and are determined to be bidirectional. The range is from 7 to 90 seconds. <p>Note This command affects fiber-optic interfaces only. Use the udld interface configuration command to enable UDLD on other interface types. For more information, see the “Enabling UDLD on an Interface” section on page 21-5.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show udld	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable UDLD globally, use the **no udld enable** global configuration command to disable normal mode UDLD on all fiber-optic ports. Use the **no udld aggressive** global configuration command to disable aggressive mode UDLD on all fiber-optic ports.

Enabling UDLD on an Interface

Beginning in privileged EXEC mode, follow these steps either to enable UDLD in the aggressive or normal mode or to disable UDLD on an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface interface-id	Enter interface configuration mode, and specify the interface to be enabled for UDLD.

	Command	Purpose
Step 3	udld port { aggressive disable }	UDLD is disabled by default. <ul style="list-style-type: none"> • udld port—Enables UDLD in normal mode on the specified interface. • udld port aggressive—Enables UDLD in aggressive mode on the specified interface. • udld port disable—Disables UDLD on the specified fiber-optic interface. This command overrides the UDLD global setting and is available only on fiber-optic interfaces. For more information about aggressive and normal modes, see the “Modes of Operation” section on page 21-1 .
Step 4	end	Return to privileged EXEC mode.
Step 5	show udld <i>interface-id</i>	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Resetting an Interface Disabled by UDLD

Beginning in privileged EXEC mode, follow these steps to reset all interfaces disabled by UDLD:

	Command	Purpose
Step 1	udld reset	Reset all interfaces disabled by UDLD.
Step 2	show udld	Verify your entries.

You can also bring up the interface by using these commands:

- The **shutdown** interface configuration command followed by the **no shutdown** interface configuration command restarts the disabled interface.
- The **no udld {aggressive | enable}** global configuration command followed by the **udld {aggressive | enable}** global configuration command re-enables the disabled interfaces.
- The **udld port disable** interface configuration command followed by the **udld port [aggressive]** interface configuration command re-enables the disabled fiber-optic interface.
- The **errdisable recovery cause udld** global configuration command enables the timer to automatically recover from the UDLD error-disabled state, and the **errdisable recovery interval interval** global configuration command specifies the time to recover from the UDLD error-disabled state.

Displaying UDLD Status

To display the UDLD status for the specified interface or for all interfaces, use the **show udld** [*interface-id*] privileged EXEC command.

For detailed information about the fields in the command output, refer to the command reference for this release.



Configuring SPAN and RSPAN

This chapter describes how to configure Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding SPAN and RSPAN, page 22-1](#)
- [Configuring SPAN and RSPAN, page 22-9](#)
- [Displaying SPAN and RSPAN Status, page 22-23](#)

Understanding SPAN and RSPAN

You can analyze network traffic passing through ports or VLANs by using SPAN or RSPAN to send a copy of the traffic to another port on the switch or on another switch that has been connected to a network analyzer or other monitoring or security device. SPAN copies (or mirrors) traffic received or sent (or both) on source ports or source VLANs to a destination port for analysis. SPAN does not affect the switching of network traffic on the source ports or VLANs. You must dedicate the destination port for SPAN use. Except for traffic that is required for the SPAN or RSPAN session, destination ports do not receive or forward traffic.

Only traffic that enters or leaves source ports or traffic that enters or leaves source VLANs can be monitored by using SPAN; traffic routed to a source VLAN cannot be monitored. For example, if incoming traffic is being monitored, traffic that gets routed from another VLAN to the source VLAN cannot be monitored; however, traffic that is received on the source VLAN and routed to another VLAN can be monitored.

You can use the SPAN or RSPAN destination port to inject traffic from a network security device. For example, if you connect a Cisco Intrusion Detection System (IDS) sensor appliance to a destination port, the IDS device can send TCP reset packets to close down the TCP session of a suspected attacker.

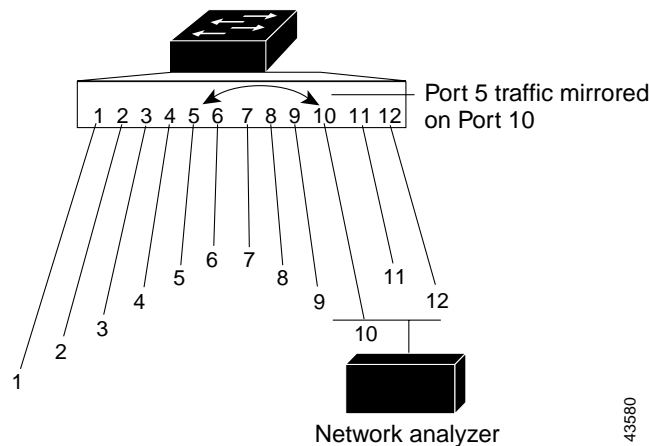
This section includes these topics:

- [Local SPAN, page 22-2](#)
- [Remote SPAN, page 22-2](#)
- [SPAN and RSPAN Concepts and Terminology, page 22-3](#)
- [SPAN and RSPAN Interaction with Other Features, page 22-8](#)

Local SPAN

Local SPAN supports a SPAN session entirely within one switch; all source ports or source VLANs and destination ports reside in the same switch. Local SPAN copies traffic from one or more source ports in any VLAN or from one or more VLANs to a destination port for analysis. For example, in [Figure 22-1](#), all traffic on port 5 (the source port) is mirrored to port 10 (the destination port). A network analyzer on port 10 receives all network traffic from port 5 without being physically attached to port 5.

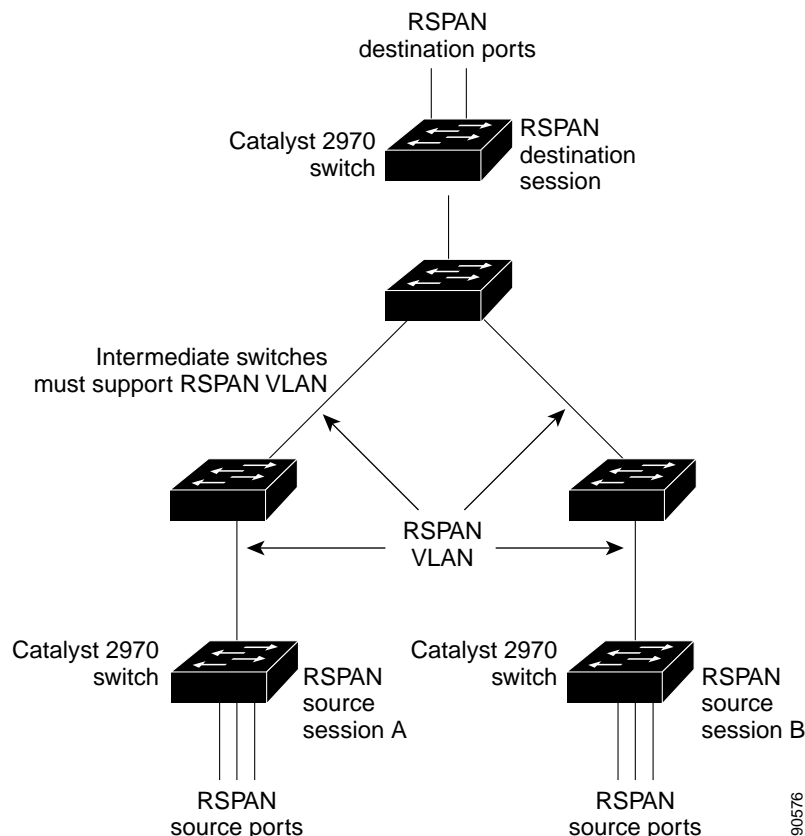
Figure 22-1 Example of Local SPAN Configuration on a Single Switch



Remote SPAN

RSPAN supports source ports, source VLANs, and destination ports on different switches, enabling remote monitoring of multiple switches across your network as shown in [Figure 22-2](#). The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The RSPAN traffic from the source ports or VLANs is copied into the RSPAN VLAN and forwarded over trunk ports carrying the RSPAN VLAN to a destination session monitoring the RSPAN VLAN. Each RSPAN source switch must have either ports or VLANs as RSPAN sources. The destination is always a physical port.

Figure 22-2 Example of RSPAN Configuration



SPAN and RSPAN Concepts and Terminology

This section describes concepts and terminology associated with SPAN and RSPAN configuration.

SPAN Sessions

SPAN sessions (local or remote) allow you to monitor traffic on one or more ports, or one or more VLANs, and send the monitored traffic to one or more destination ports.

A local SPAN session is an association of a destination port with source ports or source VLANs, all on a single network device. Local SPAN does not have separate source and destination sessions. Local SPAN sessions gather a set of ingress and egress packets specified by the user and form them into a stream of SPAN data, which is directed to the destination port.

RSPAN consists of at least one RSPAN source session, an RSPAN VLAN, and at least one RSPAN destination session. You separately configure RSPAN source sessions and RSPAN destination sessions on different network devices. To configure an RSPAN source session on a device, you associate a set of source ports or source VLANs with an RSPAN VLAN. The output of this session is the stream of SPAN packets that are sent to the RSPAN VLAN. To configure an RSPAN destination session on another device, you associate the destination port with the RSPAN VLAN. The destination session collects all RSPAN VLAN traffic and sends it out the RSPAN destination port.

An RSPAN source session is very similar to a local SPAN session, except for where the packet stream is directed. In an RSPAN source session, SPAN packets are relabeled with the RSPAN VLAN ID and directed over normal trunk ports to the destination switch.

An RSPAN destination session takes all packets received on the RSPAN VLAN, strips off the VLAN tagging, and presents them on the destination port. Its purpose is to present a copy of all RSPAN VLAN packets (except Layer 2 control packets) to the user for analysis.

There can be more than one source session and more than one destination session active in the same RSPAN VLAN. There can also be intermediate switches separating the RSPAN source and destination sessions. These switches need not be capable of running RSPAN, but they must handle the requirements of the RSPAN VLAN (see the [“RSPAN VLAN” section on page 22-8](#)).

Traffic monitoring in a SPAN session has these restrictions:

- Sources can be ports or VLANs, but you cannot mix source ports and source VLANs in the same session.
- The switch supports up to two source sessions; you can run both a local SPAN and an RSPAN source session in the same switch. The switch supports a total of 66 source and RSPAN destination sessions.
- You can have multiple destination ports in a SPAN session, but no more than 64 destination ports.
- You can configure two separate SPAN or RSPAN source sessions with separate or overlapping sets of SPAN source ports and VLANs.
- SPAN sessions do not interfere with the normal operation of the switch. However, an oversubscribed SPAN destination, for example, a 10-Mbps port monitoring a 100-Mbps port, can result in dropped or lost packets.
- When RSPAN is enabled, each packet being monitored is transmitted twice, once as normal traffic and once as a monitored packet. Therefore monitoring a large number of ports or VLANs could potentially generate large amounts of network traffic.
- You can configure SPAN sessions on disabled ports; however, a SPAN session does not become active unless you enable the destination port and at least one source port or VLAN for that session.
- The switch does not support a combination of local SPAN and RSPAN in a single session. That is, an RSPAN source session cannot have a local destination port, an RSPAN destination session cannot have a local source port, and an RSPAN destination session and an RSPAN source session that are using the same RSPAN VLAN cannot run on the same switch.

Monitored Traffic

SPAN sessions can monitor these traffic types:

- Receive (Rx) SPAN—The goal of receive (or ingress) SPAN is to monitor as much as possible all the packets received by the source interface or VLAN before any modification or processing is performed by the switch. A copy of each packet received by the source is sent to the destination port for that SPAN session.

Packets that are modified because of routing or quality of service (QoS)—for example, modified Differentiated Services Code Point (DSCP)—are copied before modification.

Features that can cause a packet to be dropped during receive processing have no effect on ingress SPAN; the destination port receives a copy of the packet even if the actual incoming packet is dropped. These features include IP standard and extended input access control lists (ACLs), ingress QoS policing, VLAN ACLs and egress QoS policing.

- **Transmit (Tx) SPAN**—The goal of transmit (or egress) SPAN is to monitor as much as possible all the packets sent by the source interface after all modification and processing is performed by the switch. A copy of each packet sent by the source is sent to the destination port for that SPAN session. The copy is provided after the packet is modified.

Features that can cause a packet to be dropped during transmit processing also affect the duplicated copy for SPAN. These features include IP standard and extended output ACLs and egress QoS policing.

- **Both**—In a SPAN session, you can also monitor a port or VLAN for both received and sent packets. This is the default.

The default configuration for local SPAN session ports is to send all packets untagged. SPAN also does not normally monitor bridge protocol data unit (BPDU) packets and Layer 2 protocols, such as Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), Dynamic Trunking Protocol (DTP), Spanning Tree Protocol (STP), and Port Aggregation Protocol (PAgP). However, when you enter the **encapsulation replicate** keywords when configuring a destination port, these changes occur:

- Packets are sent on the destination port with the same encapsulation—untagged, IEEE 802.1Q, or Inter-Switch Link (ISL)—that they had on the source port.
- Packets of all types, including BPDU and Layer 2 protocol packets are monitored.

Therefore, a local SPAN session with encapsulation replicate enabled can have a mixture of untagged, 802.1Q, and ISL tagged packets appear on the destination port.

Switch congestion can cause packets to be dropped at ingress source ports, egress source ports, or SPAN destination ports. In general, these characteristics are independent of one another. For example:

- A packet might be forwarded normally but dropped from monitoring due to an oversubscribed SPAN destination port.
- An ingress packet might be dropped from normal forwarding, but still appear on the SPAN destination port.
- An egress packet dropped because of switch congestion is also dropped from egress SPAN.

In some SPAN configurations, multiple copies of the same source packet are sent to the SPAN destination port. For example, a bidirectional (both Rx and Tx) SPAN session is configured for the Rx monitor on port A and Tx monitor on port B. If a packet enters the switch through port A and is switched to port B, both incoming and outgoing packets are sent to the destination port. Both packets are the same.

Source Ports

A source port (also called a *monitored port*) is a switched port that you monitor for network traffic analysis. In a local SPAN session or RSPAN source session, you can monitor source ports or VLANs for traffic in one or both directions. The switch supports any number of source ports (up to the maximum number of available ports on the switch) and any number of source VLANs (up to the maximum number of VLANs supported). However, the switch supports a maximum of two sessions (local or RSPAN) with source ports or VLANs and you cannot mix ports and VLANs in a single session.

A source port has these characteristics:

- It can be monitored in multiple SPAN sessions.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor.
- It can be any port type (for example, EtherChannel, Gigabit Ethernet, and so forth).
- For EtherChannel sources, you can monitor traffic for the entire EtherChannel or individually on a physical port as it participates in the port channel.

- It can be an access port, trunk port, or voice VLAN port.
- It cannot be a destination port.
- Source ports can be in the same or different VLANs.
- You can monitor multiple source ports in a single session.

Source VLANs

VLAN-based SPAN (VSPAN) is the monitoring of the network traffic in one or more VLANs. The SPAN or RSPAN source interface in VSPAN is a VLAN ID and traffic is monitored on all the ports for that VLAN.

VSPAN has these characteristics:

- All active ports in the source VLAN are included as source ports and can be monitored in either or both directions.
- On a given port, only traffic on the monitored VLAN is sent to the destination port.
- If a destination port belongs to a source VLAN, it is excluded from the source list and is not monitored.
- If ports are added to or removed from the source VLANs, the traffic on the source VLAN received by those ports is added to or removed from the sources being monitored.
- You cannot use filter VLANs in the same session with VLAN sources.
- You can monitor only Ethernet VLANs.

VLAN Filtering

When you monitor a trunk port as a source port, by default, all VLANs active on the trunk are monitored. You can limit SPAN traffic monitoring on trunk source ports to specific VLANs by using VLAN filtering.

- VLAN filtering applies only to trunk ports or to voice VLAN ports.
- VLAN filtering applies only to port-based sessions and is not allowed in sessions with VLAN sources.
- When a VLAN filter list is specified, only those VLANs in the list are monitored on trunk ports or on voice VLAN access ports.
- SPAN traffic coming from other port types is not affected by VLAN filtering; that is, all VLANs are allowed on other ports.
- VLAN filtering affects only traffic forwarded to the destination SPAN port and does not affect the switching of normal traffic.

Destination Port

Each local SPAN session or RSPAN destination session must have a destination port (also called a *monitoring port*) that receives a copy of traffic from the source ports or VLANs and sends the SPAN packets to the user, usually a network analyzer.

A destination port has these characteristics:

- For a local SPAN session, the destination port must reside on the same switch as the source port. For an RSPAN session, it is located on the switch containing the RSPAN destination session. There is no destination port on a switch running only an RSPAN source session.
- When a port is configured as a SPAN destination port, the configuration overwrites the original port configuration. When the SPAN destination configuration is removed, the port reverts to its previous configuration. If a configuration change is made to the port while it is acting as a SPAN destination port, the change does not take effect until the SPAN destination configuration had been removed.
- If the port was in an EtherChannel group, it is removed from the group while it is a destination port.
- It can be any Ethernet physical port.
- It cannot be a secure port.
- It cannot be a source port.
- It cannot be an EtherChannel group or a VLAN.
- It can participate in only one SPAN session at a time (a destination port in one SPAN session cannot be a destination port for a second SPAN session).
- When it is active, incoming traffic is disabled. The port does not transmit any traffic except that required for the SPAN session. Incoming traffic is never learned or forwarded on a destination port.
- If ingress traffic forwarding is enabled for a network security device, the destination port forwards traffic at Layer 2.
- It does not participate in any of the Layer 2 protocols (STP, VTP, CDP, DTP, PagP).
- A destination port that belongs to a source VLAN of any SPAN session is excluded from the source list and is not monitored.
- The maximum number of destination ports in a switch is 64.

Local SPAN and RSPAN destination ports behave differently regarding VLAN tagging and encapsulation:

- For local SPAN, if the **encapsulation replicate** keywords are specified for the destination port, these packets appear with the original encapsulation (untagged, ISL, or 802.1Q). If these keywords are not specified, packets appear in the untagged format. Therefore, the output of a local SPAN session with **encapsulation replicate** enabled can contain a mixture of untagged, 802.1Q, or ISL tagged packets.
- For RSPAN, the original VLAN ID is lost because it is overwritten by the RSPAN VLAN identification. Therefore, all packets appear on the destination port as untagged.

RSPAN VLAN

The RSPAN VLAN carries SPAN traffic between RSPAN source and destination sessions. It has these special characteristics:

- All traffic in the RSPAN VLAN is always flooded.
- No MAC address learning occurs on the RSPAN VLAN.
- RSPAN VLAN traffic only flows on trunk ports.
- RSPAN VLANs must be configured in VLAN configuration mode by using the **remote-span** VLAN configuration mode command.
- STP can run on RSPAN VLAN trunks but not on SPAN destination ports.

For VLANs 1 to 1005 that are visible to VLAN Trunking Protocol (VTP), the VLAN ID and its associated RSPAN characteristic are propagated by VTP. If you assign an RSPAN VLAN ID in the extended VLAN range (1006 to 4094), you must manually configure all intermediate switches.

It is normal to have multiple RSPAN VLANs in a network at the same time with each RSPAN VLAN defining a network-wide RSPAN session. That is, multiple RSPAN source sessions anywhere in the network can contribute packets to the RSPAN session. It is also possible to have multiple RSPAN destination sessions throughout the network, monitoring the same RSPAN VLAN and presenting traffic to the user. The RSPAN VLAN ID separates the sessions.

SPAN and RSPAN Interaction with Other Features

SPAN interacts with these features:

- Spanning Tree Protocol (STP)—A destination port does not participate in STP while its SPAN or RSPAN session is active. The destination port can participate in STP after the SPAN or RSPAN session is disabled. On a source port, SPAN does not affect the STP status. STP can be active on trunk ports carrying an RSPAN VLAN.
- Cisco Discovery Protocol (CDP)—A SPAN destination port does not participate in CDP while the SPAN session is active. After the SPAN session is disabled, the port again participates in CDP.
- VLAN Trunking Protocol (VTP)—You can use VTP to prune an RSPAN VLAN between switches.
- VLAN and trunking—You can modify VLAN membership or trunk settings for source or destination ports at any time. However, changes in VLAN membership or trunk settings for a destination port do not take effect until you remove the SPAN destination configuration. Changes in VLAN membership or trunk settings for a source port immediately take effect, and the respective SPAN sessions automatically adjust accordingly.
- EtherChannel—You can configure an EtherChannel group as a source port but not as a SPAN destination port. When a group is configured as a SPAN source, the entire group is monitored.

If a physical port is added to a monitored EtherChannel group, the new port is added to the SPAN source port list. If a port is removed from a monitored EtherChannel group, it is automatically removed from the source port list. If the port is the only port in the EtherChannel group, because there are no longer any ports in the group, there is no data to monitor.

A physical port that belongs to an EtherChannel group can be configured as a SPAN source port and still be a part of the EtherChannel. In this case, data from the physical port is monitored as it participates in the EtherChannel. However, if a physical port that belongs to an EtherChannel group is configured as a SPAN destination, it is removed from the group. After the port is removed from the SPAN session, it rejoins the EtherChannel group. Ports removed from an EtherChannel group remain members of the group, but they are in the *inactive* or *standalone* state.

If a physical port that belongs to an EtherChannel group is a destination port and the EtherChannel group is a source, the port is removed from the EtherChannel group and from the list of monitored ports.

- Multicast traffic can be monitored. For egress and ingress port monitoring, only a single unedited packet is sent to the SPAN destination port. It does not reflect the number of times the multicast packet is sent.
- A secure port cannot be a SPAN destination port.

For SPAN sessions, do not enable port security on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable port security on any ports with monitored egress.

- An 802.1X port can be a SPAN source port. You can enable 802.1X on a port that is a SPAN destination port; however, 802.1X is disabled until the port is removed as a SPAN destination.

For SPAN sessions, do not enable 802.1X on ports with monitored egress when ingress forwarding is enabled on the destination port. For RSPAN source sessions, do not enable 802.1X on any ports that are egress monitored.

Configuring SPAN and RSPAN

This section describes how to configure SPAN on your switch. It contains this configuration information:

- [Default SPAN and RSPAN Configuration, page 22-9](#)
- [Configuring Local SPAN, page 22-10](#)
- [Configuring RSPAN, page 22-16](#)

Default SPAN and RSPAN Configuration

[Table 22-1](#) shows the default SPAN and RSPAN configuration.

Table 22-1 Default SPAN and RSPAN Configuration

Feature	Default Setting
SPAN state (SPAN and RSPAN)	Disabled.
Source port traffic to monitor	Both received and sent traffic (both).
Encapsulation type (destination port)	Native form (untagged packets).
Ingress forwarding (destination port)	Disabled
VLAN filtering	On a trunk interface used as a source port, all VLANs are monitored.
RSPAN VLANs	None configured.

Configuring Local SPAN

This section describes how to configure Local SPAN on your switch. It contains this configuration information:

- [SPAN Configuration Guidelines, page 22-10](#)
- [Creating a Local SPAN Session, page 22-11](#)
- [Creating a Local SPAN Session and Configuring Ingress Traffic, page 22-13](#)
- [Specifying VLANs to Filter, page 22-15](#)

SPAN Configuration Guidelines

Follow these guidelines when configuring SPAN:

- You can configure a total of two local SPAN sessions or RSPAN source sessions on each switch. You can have a total of 66 SPAN sessions (local, RSPAN source, and RSPAN destination) on a switch.
- For SPAN sources, you can monitor traffic for a single port or VLAN or a series or range of ports or VLANs for each session. You cannot mix source ports and source VLANs within a single SPAN session.
- The destination port cannot be a source port; a source port cannot be a destination port.
- You cannot have two SPAN sessions using the same destination port.
- When you configure a switch port as a SPAN destination port, it is no longer a normal switch port; only monitored traffic passes through the SPAN destination port.
- Entering SPAN configuration commands does not remove previously configured SPAN parameters. You must enter the **no monitor session** {*session_number* | **all** | **local** | **remote**} global configuration command to delete configured SPAN parameters.
- For local SPAN, outgoing packets through the SPAN destination port carry the original encapsulation headers—untagged, ISL, or IEEE 802.1Q—if the **encapsulation replicate** keywords are specified. If the keywords are not specified, the packets are sent in native form. For RSPAN destination ports, outgoing packets are not tagged.
- You can configure a disabled port to be a source or destination port, but the SPAN function does not start until the destination port and at least one source port or source VLAN are enabled.
- You can limit SPAN traffic to specific VLANs by using the **filter vlan** keyword. If a trunk port is being monitored, only traffic on the VLANs specified with this keyword is monitored. By default, all VLANs are monitored on a trunk port.
- You cannot mix source VLANs and filter VLANs within a single SPAN session.

Creating a Local SPAN Session

Beginning in privileged EXEC mode, follow these steps to create a SPAN session and specify the source (monitored) ports or VLANs and the destination (monitoring) ports:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> / vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port). For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port or source VLAN to monitor. <ul style="list-style-type: none"> For source <i>interface-id</i>, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). Valid port channel numbers are 1 to 12. For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). <p>Note A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session.</p> <p>(Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen.</p> <p>(Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the SPAN monitors both sent and received traffic.</p> <ul style="list-style-type: none"> both—Monitor both received and sent traffic. This is the default. rx—Monitor received traffic. tx—Monitor sent traffic. <p>Note You can use the monitor session <i>session_number</i> source command multiple times to configure multiple source ports.</p>

	Command	Purpose
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate]}	Specify the SPAN session and the destination port (monitoring port). For <i>session_number</i> , specify the session number entered in step 3. Note For local SPAN, you must use the same session number for the source and destination interfaces. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). Note You can use monitor session <i>session_number</i> destination command multiple times to configure multiple destination ports.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a source or destination port or VLAN from the SPAN session, use the **no monitor session** *session_number* **source** {**interface** *interface-id* / **vlan** *vlan-id*} global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. For destination interfaces, the **encapsulation replicate** keywords are ignored with the **no** form of the command.

This example shows how to set up SPAN session 1 for monitoring source port traffic to a destination port. First, any existing SPAN configuration for session 1 is deleted, and then bidirectional traffic is mirrored from source Gigabit Ethernet port 0/1 to destination Gigabit Ethernet port 0/3, retaining the encapsulation method.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/1
Switch(config)# monitor session 1 destination interface gigabitethernet0/3
encapsulation replicate
Switch(config)# end
```

This example shows how to remove port 0/1 as a SPAN source for SPAN session 1:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1
Switch(config)# end
```

This example shows how to disable received traffic monitoring on port 0/1, which was configured for bidirectional monitoring:

```
Switch(config)# no monitor session 1 source interface gigabitethernet0/1 rx
```

The monitoring of traffic received on port 0/1 is disabled, but traffic sent from this port continues to be monitored.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on all ports belonging to VLANs 1 through 3, and send it to destination Gigabit Ethernet port 0/2. The configuration is then modified to also monitor all traffic on all ports belonging to VLAN 10.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source vlan 1 - 3 rx
Switch(config)# monitor session 2 destination interface gigabitethernet0/2
Switch(config)# monitor session 2 source vlan 10
Switch(config)# end
```

Creating a Local SPAN Session and Configuring Ingress Traffic

Beginning in privileged EXEC mode, follow these steps to create a SPAN session, to specify the source ports or VLANs and the destination ports, and to enable ingress traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).



Note

Refer to the [“Creating a Local SPAN Session” section on page 22-11](#) for details about the keywords not related to ingress traffic.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the SPAN session and the source port (monitored port).

	Command	Purpose
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate] [ingress { dot1q vlan <i>vlan-id</i> isl untagged vlan <i>vlan-id</i> vlan <i>vlan-id</i> }]}	Specify the SPAN session, the destination port, the packet encapsulation, and the ingress VLAN and encapsulation. For <i>session_number</i> , specify the session number entered in step 3. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma or hyphen. (Optional) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged). Enter ingress with keywords to enable ingress traffic forwarding on the destination port and specify the encapsulation type: <ul style="list-style-type: none"> • dot1q vlan <i>vlan-id</i>—Forward ingress packets with 802.1Q encapsulation with the specified VLAN as the default VLAN. • isl—Forward ingress packets with ISL encapsulation. • untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Forward ingress packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a source or destination port or VLAN from the SPAN session, use the **no monitor session** *session_number* **source** {**interface** *interface-id* | **vlan** *vlan-id*} global configuration command or the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. For destination interfaces, the encapsulation and ingress options are ignored with the **no** form of the command.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor received traffic on all source Gigabit Ethernet source port 0/1, and send it to destination Gigabit Ethernet port 0/2 with the same egress encapsulation type as the source port, and to enable ingress forwarding with 802.1Q encapsulation and VLAN 6 as the default ingress VLAN.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source gigabitethernet0/1 rx
Switch(config)# monitor session 2 destination interface gigabitethernet0/2 encapsulation
replicate ingress dot1q vlan 6
Switch(config)# end
```

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to limit SPAN source traffic to specific VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source interface <i>interface-id</i>	Specify the characteristics of the source port (monitored port) and SPAN session. For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	monitor session <i>session_number</i> filter vlan <i>vlan-id</i> [, -]	Limit the SPAN source traffic to specific VLANs. For <i>session_number</i> , enter the session number specified in Step 3. For <i>vlan-id</i> , the range is 1 to 4094. (Optional) Use a comma (,) to specify a series of VLANs, or use a hyphen (-) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [encapsulation replicate] }	Specify the SPAN session and the destination port (monitoring port). For <i>session_number</i> , specify the session number entered in step 3. For <i>interface-id</i> , specify the destination port. The destination interface must be a physical port; it cannot be an EtherChannel, and it cannot be a VLAN. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. (Optional) Enter encapsulation replicate to specify that the destination interface replicates the source interface encapsulation method. If not selected, the default is to send packets in native form (untagged).
Step 6	end	Return to privileged EXEC mode.
Step 7	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter** global configuration command.

This example shows how to remove any existing configuration on SPAN session 2, configure SPAN session 2 to monitor traffic received on Gigabit Ethernet trunk port 0/4, and send traffic for only VLANs 1 through 5 and 9 to destination Gigabit Ethernet port 0/3.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet0/4 rx
Switch(config)# monitor session 2 filter vlan 1 - 5 , 9
Switch(config)# monitor session 2 destination interface gigabitethernet0/3
Switch(config)# end
```

Configuring RSPAN

This section describes how to configure RSPAN on your switch. It contains this configuration information:

- [RSPAN Configuration Guidelines, page 22-16](#)
- [Configuring a VLAN as an RSPAN VLAN, page 22-17](#)
- [Creating an RSPAN Source Session, page 22-18](#)
- [Creating an RSPAN Destination Session, page 22-19](#)
- [Creating an RSPAN Destination Session, page 22-19](#)
- [Specifying VLANs to Filter, page 22-22](#)

RSPAN Configuration Guidelines

Follow these guidelines when configuring RSPAN:

- All the items in the [“SPAN Configuration Guidelines” section on page 22-10](#) apply to RSPAN.
- As RSPAN VLANs have special properties, you should reserve a few VLANs across your network for use as RSPAN VLANs; do not assign access ports to these VLANs.
- You can apply an output access control list (ACL) to RSPAN traffic to selectively filter or monitor specific packets. Specify these ACLs on the RSPAN VLAN in the RSPAN source switches.
- For RSPAN configuration, you can distribute the source ports and the destination ports across multiple switches in your network.
- RSPAN does not support BPDU packet monitoring or other Layer 2 switch protocols.
- The RSPAN VLAN is configured only on trunk ports and not on access ports. To avoid unwanted traffic in RSPAN VLANs, make sure that the VLAN remote-span feature is supported in all the participating switches.
- Access ports (including voice VLAN ports) on the RSPAN VLAN are put in the inactive state.
- RSPAN VLANs are included as sources for port-based RSPAN sessions when source trunk ports have active RSPAN VLANs. RSPAN VLANs can also be sources in SPAN sessions. However, since the switch does not monitor spanned traffic, it does not support egress spanning of packets on any RSPAN VLAN identified as the destination of an RSPAN source session on the switch.
- You can configure any VLAN as an RSPAN VLAN as long as these conditions are met:
 - The same RSPAN VLAN is used for an RSPAN session in all the switches.
 - All participating switches support RSPAN.

- We recommend that you configure an RSPAN VLAN before you configure an RSPAN source or a destination session.
- If you enable VTP and VTP pruning, RSPAN traffic is pruned in the trunks to prevent the unwanted flooding of RSPAN traffic across the network for VLAN IDs that are lower than 1005.

Configuring a VLAN as an RSPAN VLAN

First create a new VLAN to be the RSPAN VLAN for the RSPAN session. You must create the RSPAN VLAN in all switches that will participate in RSPAN. If the RSPAN VLAN-ID is in the normal range (lower than 1005) and VTP is enabled in the network, you can create the RSPAN VLAN in one switch, and VTP propagates it to the other switches in the VTP domain. For extended-range VLANs (greater than 1005), you must configure RSPAN VLAN on both source and destination switches and any intermediate switches.

Use VTP pruning to get an efficient flow of RSPAN traffic, or manually delete the RSPAN VLAN from all trunks that do not need to carry the RSPAN traffic.

Beginning in privileged EXEC mode, follow these steps to create an RSPAN VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan <i>vlan-id</i>	Enter a VLAN ID to create a VLAN, or enter the VLAN ID of an existing VLAN, and enter VLAN configuration mode. The range is from 2 to 1001 and from 1006 to 4094. Note The RSPAN VLAN cannot be VLAN 1 (the default VLAN) or VLAN IDs 1002 through 1005 (reserved for Token Ring and FDDI VLANs).
Step 3	remote-span	Configure the VLAN as an RSPAN VLAN.
Step 4	end	Return to privileged EXEC mode.
Step 5	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To remove the remote SPAN characteristic from a VLAN and convert it back to a normal VLAN, use the **no remote-span** VLAN configuration command.

This example shows how to create RSPAN VLAN 901.

```
Switch(config)# vlan 901
Switch(config-vlan)# remote span
Switch(config-vlan)# end
```

Creating an RSPAN Source Session

Beginning in privileged EXEC mode, follow these steps to start an RSPAN source session and to specify the monitored source and the destination RSPAN VLAN:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing RSPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source { interface <i>interface-id</i> / vlan <i>vlan-id</i> } [, -] [both rx tx]	Specify the RSPAN session and the source port (monitored port). For <i>session_number</i> , the range is from 1 to 66. Enter a source port or source VLAN for the RSPAN session: <ul style="list-style-type: none"> For <i>interface-id</i>, specify the source port to monitor. Valid interfaces include physical interfaces and port-channel logical interfaces (port-channel <i>port-channel-number</i>). Valid port channel numbers are 1 to 12. For <i>vlan-id</i>, specify the source VLAN to monitor. The range is 1 to 4094 (excluding the RSPAN VLAN). <p>Note A single session can include multiple sources (ports or VLANs), defined in a series of commands, but you cannot combine source ports and source VLANs in one session.</p> <p>(Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen.</p> <p>(Optional) Specify the direction of traffic to monitor. If you do not specify a traffic direction, the source interface sends both sent and received traffic.</p> <ul style="list-style-type: none"> both—Monitor both received and sent traffic. rx—Monitor received traffic. tx—Monitor sent traffic.
Step 4	monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i>	Specify the RSPAN session and the destination RSPAN VLAN. For <i>session_number</i> , enter the number defined in Step 3. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command.

To remove a source port or VLAN from the SPAN session, use the **no monitor session *session_number* source {*interface interface-id* | *vlan vlan-id*}** global configuration command. To remove the RSPAN VLAN from the session, use the **no monitor session *session_number* destination remote *vlan vlan-id***.

This example shows how to remove any existing RSPAN configuration for session 1, configure RSPAN session 1 to monitor multiple source interfaces, and configure the destination as RSPAN VLAN 901.

```
Switch(config)# no monitor session 1
Switch(config)# monitor session 1 source interface gigabitethernet0/1 tx
Switch(config)# monitor session 1 source interface gigabitethernet0/2 rx
Switch(config)# monitor session 1 source interface gigabitethernet0/3
Switch(config)# monitor session 1 source interface port-channel 12
Switch(config)# monitor session 1 destination remote vlan 901
Switch(config)# end
```

Creating an RSPAN Destination Session

You configure the RSPAN destination session on a different switch; that is, not the switch on which the source session was configured.

Beginning in privileged EXEC mode, follow these steps to define the RSPAN VLAN on that switch, to create an RSPAN destination session, and to specify the source RSPAN VLAN and the destination port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan <i>vlan-id</i>	Enter the VLAN ID of the RSPAN VLAN created from the source switch, and enter VLAN configuration mode. Note If both switches are participating in VTP and the RSPAN VLAN ID is from 2 to 1005, Steps 2 through 4 are not required because the RSPAN VLAN ID is propagated through the VTP network.
Step 3	remote-span	Identify the VLAN as the RSPAN VLAN.
Step 4	exit	Return to global configuration mode.
Step 5	no monitor session {<i>session_number</i> all local remote}	Remove any existing RSPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all RSPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 6	monitor session <i>session_number</i> source remote <i>vlan vlan-id</i>	Specify the RSPAN session and the source RSPAN VLAN. For <i>session_number</i> , the range is from 1 to 66. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.

	Command	Purpose
Step 7	monitor session <i>session_number</i> destination interface <i>interface-id</i>	Specify the RSPAN session and the destination interface. For <i>session_number</i> , enter the number defined in Step 6. Note In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. For <i>interface-id</i> , specify the destination interface. The destination interface must be a physical interface. Note Though visible in the command-line help string, encapsulation replicate is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged.
Step 8	end	Return to privileged EXEC mode.
Step 9	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete a SPAN session, use the **no monitor session** *session_number* global configuration command. To remove a destination port from the SPAN session, use the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. To remove the RSPAN VLAN from the session, use the **no monitor session** *session_number* **source remote vlan** *vlan-id*.

This example shows how to configure VLAN 901 as the source remote VLAN and port 0/4 as the destination interface:

```
Switch(config)# monitor session 1 source remote vlan 901
Switch(config)# monitor session 1 destination interface gigabitethernet0/4
Switch(config)# end
```

Creating an RSPAN Destination Session and Configuring Ingress Traffic

Beginning in privileged EXEC mode, follow these steps to create an RSPAN destination session, to specify the source RSPAN VLAN and the destination port, and to enable ingress traffic on the destination port for a network security device (such as a Cisco IDS Sensor Appliance).



Note

Refer to the [“Creating an RSPAN Destination Session” section on page 22-19](#) for details about the keywords not related to ingress traffic. This procedure assumes the RSPAN VLAN has already been configured.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session.

	Command	Purpose
Step 3	monitor session <i>session_number</i> source remote vlan <i>vlan-id</i>	Specify the RSPAN session and the source RSPAN VLAN. For <i>session_number</i> , the range is from 1 to 66. For <i>vlan-id</i> , specify the source RSPAN VLAN to monitor.
Step 4	monitor session <i>session_number</i> destination { interface <i>interface-id</i> [, -] [ingress { dot1q vlan <i>vlan-id</i> isl untagged vlan <i>vlan-id</i> vlan <i>vlan-id</i> }]}	Specify the SPAN session, the destination port, the packet encapsulation, and the ingress VLAN and encapsulation. For <i>session_number</i> , enter the number defined in Step 4. Note In an RSPAN destination session, you must use the same session number for the source RSPAN VLAN and the destination port. For <i>interface-id</i> , specify the destination interface. The destination interface must be a physical interface. Note Though visible in the command-line help string, encapsulation replicate is not supported for RSPAN. The original VLAN ID is overwritten by the RSPAN VLAN ID, and all packets appear on the destination port as untagged. (Optional) [, -] Specify a series or range of interfaces. Enter a space before and after the comma; enter a space before and after the hyphen. Enter ingress with additional keywords to enable ingress traffic forwarding on the destination port and to specify the encapsulation type: <ul style="list-style-type: none"> • dot1q vlan <i>vlan-id</i>—Forward ingress packets with 802.1Q encapsulation with the specified VLAN as the default VLAN. • isl—Forward ingress packets with ISL encapsulation. • untagged vlan <i>vlan-id</i> or vlan <i>vlan-id</i>—Forward ingress packets with untagged encapsulation type with the specified VLAN as the default VLAN.
Step 5	end	Return to privileged EXEC mode.
Step 6	show monitor [session <i>session_number</i>] show running-config	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To delete an RSPAN session, use the **no monitor session** *session_number* global configuration command. To remove a destination port from the RSPAN session, use the **no monitor session** *session_number* **destination interface** *interface-id* global configuration command. The ingress options are ignored with the **no** form of the command.

This example shows how to configure VLAN 901 as the source remote VLAN in RSPAN session 2, to configure Gigabit Ethernet source port 0/2 as the destination interface, and to enable ingress forwarding on the interface with VLAN 6 as the default ingress VLAN.

```
Switch(config)# monitor session 2 source remote vlan 901
Switch(config)# monitor session 2 destination interface gigabitethernet0/2 ingress vlan 6
Switch(config)# end
```

Specifying VLANs to Filter

Beginning in privileged EXEC mode, follow these steps to configure the RSPAN source session to limit RSPAN source traffic to specific VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no monitor session { <i>session_number</i> all local remote }	Remove any existing SPAN configuration for the session. For <i>session_number</i> , the range is from 1 to 66. Specify all to remove all SPAN sessions, local to remove all local sessions, or remote to remove all remote SPAN sessions.
Step 3	monitor session <i>session_number</i> source interface <i>interface-id</i>	Specify the characteristics of the source port (monitored port) and SPAN session. For <i>session_number</i> , the range is from 1 to 66. For <i>interface-id</i> , specify the source port to monitor. The interface specified must already be configured as a trunk port.
Step 4	monitor session <i>session_number</i> filter vlan <i>vlan-id</i> [, -]	Limit the SPAN source traffic to specific VLANs. For <i>session_number</i> , enter the session number specified in step 3. For <i>vlan-id</i> , the range is 1 to 4094. (Optional) Use a comma (,) to specify a series of VLANs or use a hyphen (-) to specify a range of VLANs. Enter a space before and after the comma; enter a space before and after the hyphen.
Step 5	monitor session <i>session_number</i> destination remote vlan <i>vlan-id</i>	Specify the RSPAN session and the destination remote VLAN (RSPAN VLAN). For <i>session_number</i> , enter the session number specified in step 3. For <i>vlan-id</i> , specify the RSPAN VLAN to carry the monitored traffic to the destination port.
Step 6	end	Return to privileged EXEC mode.
Step 7	show monitor [<i>session session_number</i>] show running-config	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save the configuration in the configuration file.

To monitor all VLANs on the trunk port, use the **no monitor session** *session_number* **filter vlan** global configuration command.

This example shows how to remove any existing configuration on RSPAN session 2, configure RSPAN session 2 to monitor traffic received on trunk port 0/4, and send traffic for only VLANs 1 through 5 and 9 to destination RSPAN VLAN 902.

```
Switch(config)# no monitor session 2
Switch(config)# monitor session 2 source interface gigabitethernet0/4 rx
Switch(config)# monitor session 2 filter vlan 1 - 5 , 9
Switch(config)# monitor session 2 destination remote vlan 902
Switch(config)# end
```

Displaying SPAN and RSPAN Status

To display the current SPAN or RSPAN configuration, use the **show monitor** user EXEC command. You can also use the **show running-config** privileged EXEC command to display configured SPAN or RSPAN sessions.

This is an example of output for a switch with two source sessions configured:

```
Switch# show monitor
Session 1
-----
Type           :Local Session
Source Ports:
  RX Only:      Gi0/24
  TX Only:      None
  Both:         Gi0/1-2,Gi0/5-6
Source VLANs:
  RX Only:      None
  TX Only:      None
  Both:         None
Source RSPAN VLAN:None
Destination Ports:Gi0/18
  Encapsulation:Replicate
Filter VLANs:   None
Dest RSPAN VLAN: None

Session 2
-----
Type           :Remote Source Session
Source Ports:
  RX Only:      None
  TX Only:      None
  Both:         None
Source VLANs:
  RX Only:      None
  TX Only:      10
  Both:         1-9
Source RSPAN VLAN:None
Destination Ports:None
Filter VLANs:   None
Dest RSPAN VLAN: 105
```

This is an example of output for the **show monitor session all** user EXEC command when ingress traffic forwarding is enabled:

```
Switch# show monitor session all
Session 1
-----
Type                :Local Session
Source Ports        :
    Both            :Gi0/2
Destination Ports   :Gi0/3
Encapsulation       :Replicate
    Ingress:Enabled, default VLAN = 5
    Ingress encapsulation:DOT1Q

Session 2
-----
Type                :Local Session
Source Ports        :
    Both            :Gi0/1
Destination Ports   :Gi0/4
Encapsulation       :Replicate
    Ingress:Enabled
    Ingress encapsulation:ISL
```

This is an example of the configuration and output for the **show running-config** privileged EXEC command when ingress traffic forwarding is enabled. SPAN and RSPAN sessions are displayed near the end of the output.

```
Switch(config)# monitor session 2 source remote vlan 901
Switch(config)# monitor session 2 destination interface gigabitethernet0/2 ingress vlan 6
Switch(config)# end
Switch# show running-config
Building configuration...

Current configuration : 8238 bytes
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log datetime
no service password-encryption
service sequence-numbers

<output truncated>

!
!
monitor session 2 destination interface Gi0/2 ingress vlan 6
end
```



Configuring RMON

This chapter describes how to configure Remote Network Monitoring (RMON) on the Catalyst 2970 switch.

RMON is a standard monitoring specification that defines a set of statistics and functions that can be exchanged between RMON-compliant console systems and network probes. RMON provides you with comprehensive network-fault diagnosis, planning, and performance-tuning information.



Note

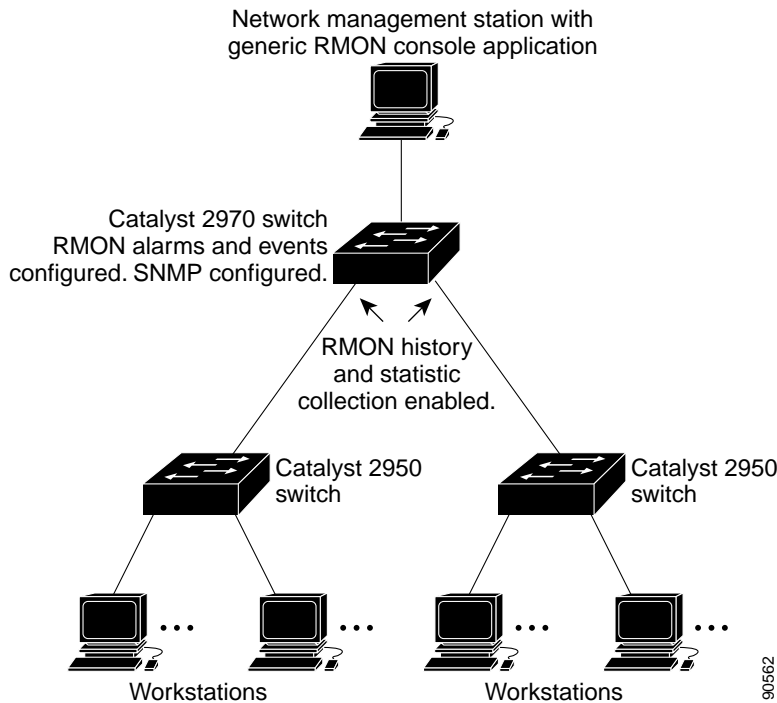
For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding RMON, page 23-1](#)
- [Configuring RMON, page 23-2](#)
- [Displaying RMON Status, page 23-6](#)

Understanding RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON feature with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments as shown in [Figure 23-1](#).

Figure 23-1 Remote Monitoring Example

The switch supports these RMON groups (defined in RFC 1757):

- **Statistics (RMON group 1)**—Collects Ethernet statistics (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) on an interface.
- **History (RMON group 2)**—Collects a history group of statistics on Ethernet interfaces (including Fast Ethernet and Gigabit Ethernet statistics, depending on the switch type and supported interfaces) for a specified polling interval.
- **Alarm (RMON group 3)**—Monitors a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms can be used with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- **Event (RMON group 9)**—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry or an SNMP trap.

Because switches supported by this software release use hardware counters for RMON data processing, the monitoring is more efficient, and little processing power is required.

Configuring RMON

These sections describe how to configure RMON on your switch:

- [Default RMON Configuration, page 23-3](#)
- [Configuring RMON Alarms and Events, page 23-3](#) (required)
- [Collecting Group History Statistics on an Interface, page 23-5](#) (optional)
- [Collecting Group Ethernet Statistics on an Interface, page 23-6](#) (optional)

Default RMON Configuration

RMON is disabled by default; no alarms or events are configured.

Only RMON 1 is supported on the switch.

Configuring RMON Alarms and Events

You can configure your switch for RMON by using the command-line interface (CLI) or an SNMP-compatible network management station. We recommend that you use a generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. You must also configure SNMP on the switch to access RMON MIB objects. For more information, see [Chapter 25, “Configuring SNMP.”](#)

Beginning in privileged EXEC mode, follow these steps to enable RMON alarms and events. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	rmon alarm <i>number variable interval</i> { absolute delta } rising-threshold <i>value</i> [<i>event-number</i>] falling-threshold <i>value</i> [<i>event-number</i>] [owner <i>string</i>]	Set an alarm on a MIB object. <ul style="list-style-type: none"> For <i>number</i>, specify the alarm number. The range is 1 to 65535. For <i>variable</i>, specify the MIB object to monitor. For <i>interval</i>, specify the time in seconds the alarm monitors the MIB variable. The range is 1 to 4294967295 seconds. Specify the absolute keyword to test each MIB variable directly. Specify the delta keyword to test the change between samples of a MIB variable. For <i>value</i>, specify a number at which the alarm is triggered and one for when the alarm is reset. The range for the rising threshold and falling threshold values is -2147483648 to 2147483647. (Optional) For <i>event-number</i>, specify the event number to trigger when the rising or falling threshold exceeds its limit. (Optional) For owner <i>string</i>, specify the owner of the alarm.

	Command	Purpose
Step 3	rmon event <i>number</i> [description <i>string</i>] [log] [owner <i>string</i>] [trap <i>community</i>]	Add an event in the RMON event table that is associated with an RMON event number. <ul style="list-style-type: none"> For <i>number</i>, assign an event number. The range is 1 to 65535. (Optional) For description <i>string</i>, specify a description of the event. (Optional) Use the log keyword to generate an RMON log entry when the event is triggered. (Optional) For owner <i>string</i>, specify the owner of this event. (Optional) For trap <i>community</i>, enter the SNMP community string used for this trap.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an alarm, use the **no rmon alarm** *number* global configuration command on each alarm you configured. You cannot disable at once all the alarms that you configured. To disable an event, use the **no rmon event** *number* global configuration command. To learn more about alarms and events and how they interact with each other, refer to RFC 1757.

You can set an alarm on any MIB object. The following example configures RMON alarm number 10 by using the **rmon alarm** command. The alarm monitors the MIB variable *ifEntry.20.1* once every 20 seconds until the alarm is disabled and checks the change in the variable's rise or fall. If the *ifEntry.20.1* value shows a MIB counter increase of 15 or more, such as from 100000 to 100015, the alarm is triggered. The alarm in turn triggers event number 1, which is configured with the **rmon event** command. Possible events can include a log entry or an SNMP trap. If the *ifEntry.20.1* value changes by 0, the alarm is reset and can be triggered again.

```
Switch(config)# rmon alarm 10 ifEntry.20.1 20 delta rising-threshold 15 1
falling-threshold 0 owner jjohnson
```

The following example creates RMON event number 1 by using the **rmon event** command. The event is defined as *High ifOutErrors* and generates a log entry when the event is triggered by the alarm. The user *jjones* owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.

```
Switch(config)# rmon event 1 log trap eventtrap description "High ifOutErrors" owner
jjones
```

Collecting Group History Statistics on an Interface

You must first configure RMON alarms and events to display collection information.

Beginning in privileged EXEC mode, follow these steps to collect group history statistics on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to collect history.
Step 3	rmon collection history <i>index</i> [buckets <i>bucket-number</i>] [interval <i>seconds</i>] [owner <i>ownername</i>]	Enable history collection for the specified number of buckets and time period. <ul style="list-style-type: none"> For <i>index</i>, identify the RMON group of statistics. The range is 1 to 65535. (Optional) For buckets <i>bucket-number</i>, specify the maximum number of buckets desired for the RMON collection history group of statistics. The range is 1 to 65535. The default is 50 buckets. (Optional) For interval <i>seconds</i>, specify the number of seconds in each polling cycle. The range is 1 to 3600. The default is 1800 seconds. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	show rmon history	Display the contents of the switch history table.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable history collection, use the **no rmon collection history** *index* interface configuration command.

Collecting Group Ethernet Statistics on an Interface

Beginning in privileged EXEC mode, follow these steps to collect group Ethernet statistics on an interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface on which to collect statistics.
Step 3	rmon collection stats <i>index</i> [owner <i>ownername</i>]	Enable RMON statistic collection on the interface. <ul style="list-style-type: none"> For <i>index</i>, specify the RMON group of statistics. The range is from 1 to 65535. (Optional) For owner <i>ownername</i>, enter the name of the owner of the RMON group of statistics.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	show rmon statistics	Display the contents of the switch statistics table.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the collection of group Ethernet statistics, use the **no rmon collection stats** *index* interface configuration command.

This example shows how to collect RMON statistics for the owner *root* on Gigabit Ethernet interface 0/1:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# rmon collection stats 2 owner root
```

Displaying RMON Status

To display the RMON status, use one or more of the privileged EXEC commands in [Table 23-1](#):

Table 23-1 Commands for Displaying RMON Status

Command	Purpose
show rmon	Displays general RMON statistics.
show rmon alarms	Displays the RMON alarm table.
show rmon events	Displays the RMON event table.
show rmon history	Displays the RMON history table.
show rmon statistics	Displays the RMON statistics table.

For information about the fields in these displays, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.



Configuring System Message Logging

This chapter describes how to configure system message logging on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding System Message Logging, page 24-1](#)
- [Configuring System Message Logging, page 24-2](#)
- [Displaying the Logging Configuration, page 24-12](#)

Understanding System Message Logging

By default, a switch sends the output from system messages and **debug** privileged EXEC commands to a logging process. The logging process controls the distribution of logging messages to various destinations, such as the logging buffer, terminal lines, or a UNIX syslog server, depending on your configuration. The process also sends messages to the console.



Note

The syslog format is compatible with 4.3 BSD UNIX.

When the logging process is disabled, messages are sent only to the console. The messages are sent as they are generated, so message and debug output are interspersed with prompts or output from other commands. Messages appear on the console after the process that generated them has finished.

You can set the severity level of the messages to control the type of messages displayed on the consoles and each of the destinations. You can time stamp log messages or set the syslog source address to enhance real-time debugging and management. For information on possible messages, refer to the system message guide for this release.

You can access logged system messages by using the switch command-line interface (CLI) or by saving them to a properly configured syslog server. The switch software saves syslog messages in an internal buffer.

You can remotely monitor system messages by viewing the logs on a syslog server or by accessing the switch through Telnet or through the console port.

Configuring System Message Logging

These sections describe how to configure system message logging:

- [System Log Message Format, page 24-2](#)
- [Default System Message Logging Configuration, page 24-3](#)
- [Disabling Message Logging, page 24-4](#) (optional)
- [Setting the Message Display Destination Device, page 24-4](#) (optional)
- [Synchronizing Log Messages, page 24-5](#) (optional)
- [Enabling and Disabling Time Stamps on Log Messages, page 24-7](#) (optional)
- [Enabling and Disabling Sequence Numbers in Log Messages, page 24-7](#) (optional)
- [Defining the Message Severity Level, page 24-8](#) (optional)
- [Limiting Syslog Messages Sent to the History Table and to SNMP, page 24-9](#) (optional)
- [Configuring UNIX Syslog Servers, page 24-10](#) (optional)

System Log Message Format

System log messages can contain up to 80 characters and a percent sign (%), which follows the optional sequence number or time-stamp information, if configured. Messages appear in this format:

seq no:timestamp: %facility-severity-MNEMONIC:description

The part of the message preceding the percent sign depends on the setting of the **service sequence-numbers**, **service timestamps log datetime**, **service timestamps log datetime [localtime] [msec] [show-timezone]**, or **service timestamps log uptime** global configuration command.

[Table 24-1](#) describes the elements of syslog messages.

Table 24-1 System Log Message Elements

Element	Description
<i>seq no:</i>	Stamps log messages with a sequence number only if the service sequence-numbers global configuration command is configured. For more information, see the “ Enabling and Disabling Sequence Numbers in Log Messages ” section on page 24-7.
<i>timestamp</i> formats: <i>mm/dd hh:mm:ss</i> or <i>hh:mm:ss</i> (short uptime) or <i>d h</i> (long uptime)	Date and time of the message or event. This information appears only if the service timestamps log [datetime log] global configuration command is configured. For more information, see the “ Enabling and Disabling Time Stamps on Log Messages ” section on page 24-7.
<i>facility</i>	The facility to which the message refers (for example, SNMP, SYS, and so forth). For a list of supported facilities, see Table 24-4 on page 24-12 .
<i>severity</i>	Single-digit code from 0 to 7 that is the severity of the message. For a description of the severity levels, see Table 24-3 on page 24-9 .

Table 24-1 System Log Message Elements (continued)

Element	Description
<i>MNEMONIC</i>	Text string that uniquely describes the message.
<i>description</i>	Text string containing detailed information about the event being reported.

This example shows a partial switch system message:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/1, changed state to up
00:00:47: %LINK-3-UPDOWN: Interface GigabitEthernet0/2, changed state to up
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan1, changed state to down
00:00:48: %LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet0/1, changed
state to down 2
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
18:47:02: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
*Mar  1 18:48:50.483 UTC: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Default System Message Logging Configuration

[Table 24-2](#) shows the default system message logging configuration.

Table 24-2 Default System Message Logging Configuration

Feature	Default Setting
System message logging to the console	Enabled.
Console severity	Debugging (and numerically lower levels; see Table 24-3 on page 24-9).
Logging file configuration	No filename specified.
Logging buffer size	4096 bytes.
Logging history size	1 message.
Time stamps	Disabled.
Synchronous logging	Disabled.
Logging server	Disabled.
Syslog server IP address	None configured.
Server facility	Local7 (see Table 24-4 on page 24-12).
Server severity	Informational (and numerically lower levels; see Table 24-3 on page 24-9).

Disabling Message Logging

Message logging is enabled by default. It must be enabled to send messages to any destination other than the console. When enabled, log messages are sent to a logging process, which logs messages to designated locations asynchronously to the processes that generated the messages.

Beginning in privileged EXEC mode, follow these steps to disable message logging. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no logging on	Disable message logging.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config or show logging	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Disabling the logging process can slow down the switch because a process must wait until the messages are written to the console before continuing. When the logging process is disabled, messages appear on the console as soon as they are produced, often appearing in the middle of command output.

The **logging synchronous** global configuration command also affects the display of messages to the console. When this command is enabled, messages appear only after you press Return. For more information, see the [“Synchronizing Log Messages” section on page 24-5](#).

To re-enable message logging after it has been disabled, use the **logging on** global configuration command.

Setting the Message Display Destination Device

If message logging is enabled, you can send messages to specific locations in addition to the console. Beginning in privileged EXEC mode, use one or more of the following commands to specify the locations that receive messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging buffered <i>[size]</i>	Log messages to an internal buffer on the switch. The default buffer size is 4096. The range is 4096 to 2147483647 bytes. If the switch fails, the log file is lost unless you previously saved it to Flash memory. See Step 4. Note Do not make the buffer size too large because the switch could run out of memory for other tasks. Use the show memory privileged EXEC command to view the free processor memory on the switch. However, this value is the maximum available, and the buffer size should <i>not</i> be set to this amount.

	Command	Purpose
Step 3	logging <i>host</i>	Log messages to a UNIX syslog server host. For <i>host</i> , specify the name or IP address of the host to be used as the syslog server. To build a list of syslog servers that receive logging messages, enter this command more than once. For complete syslog server configuration steps, see the “Configuring UNIX Syslog Servers” section on page 24-10.
Step 4	logging file flash : <i>filename</i> [<i>max-file-size</i> [<i>min-file-size</i>]] [<i>severity-level-number</i> <i>type</i>]	Store log messages in a file in Flash memory. <ul style="list-style-type: none">For <i>filename</i>, enter the log message filename.(Optional) For <i>max-file-size</i>, specify the maximum logging file size. The range is 4096 to 2147483647. The default is 4096 bytes.(Optional) For <i>min-file-size</i>, specify the minimum logging file size. The range is 1024 to 2147483647. The default is 2048 bytes.(Optional) For <i>severity-level-number</i> <i>type</i>, specify either the logging severity level or the logging type. The severity range is 0 to 7. For a list of logging type keywords, see Table 24-3 on page 24-9. By default, the log file receives debugging messages and numerically lower levels.
Step 5	end	Return to privileged EXEC mode.
Step 6	terminal monitor	Log messages to a nonconsole terminal during the current session. Terminal parameter-setting commands are set locally and do not remain in effect after the session has ended. You must perform this step for each session to see the debugging messages.
Step 7	show running-config	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **logging buffered** global configuration command copies logging messages to an internal buffer. The buffer is circular, so newer messages overwrite older messages after the buffer is full. To display the messages that are logged in the buffer, use the **show logging** privileged EXEC command. The first message displayed is the oldest message in the buffer. To clear the contents of the buffer, use the **clear logging** privileged EXEC command.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a file, use the **no logging file** [*severity-level-number* | *type*] global configuration command.

Synchronizing Log Messages

You can synchronize unsolicited messages and **debug** privileged EXEC command output with solicited device output and prompts for a specific console port line or virtual terminal line. You can identify the types of messages to be output asynchronously based on the level of severity. You can also determine the maximum number of buffers for storing asynchronous messages for the terminal after which messages are dropped.

When synchronous logging of unsolicited messages and **debug** command output is enabled, unsolicited device output appears on the console or printed after solicited device output appears or printed.

Unsolicited messages and **debug** command output appears on the console after the prompt for user input

is returned. Therefore, unsolicited messages and **debug** command output are not interspersed with solicited device output and prompts. After the unsolicited messages appear, the console again displays the user prompt.

Beginning in privileged EXEC mode, follow these steps to configure synchronous logging. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line [console vty] <i>line-number</i> [<i>ending-line-number</i>]	Specify the line to be configured for synchronous logging of messages. <ul style="list-style-type: none"> Use the console keyword for configurations that occur through the switch console port. Use the line vty line-number command to specify which vty lines are to have synchronous logging enabled. You use a vty connection for configurations that occur through a Telnet session. The range of line numbers is from 0 to 15. <p>You can change the setting of all 16 vty lines at once by entering: line vty 0 15</p> <p>Or you can change the setting of the single vty line being used for your current connection. For example, to change the setting for vty line 2, enter: line vty 2</p> <p>When you enter this command, the mode changes to line configuration.</p>
Step 3	logging synchronous [level <i>severity-level</i> all] limit <i>number-of-buffers</i>	Enable synchronous logging of messages. <ul style="list-style-type: none"> (Optional) For level severity-level, specify the message severity level. Messages with a severity level equal to or higher than this value are printed asynchronously. Low numbers mean greater severity and high numbers mean lesser severity. The default is 2. (Optional) Specifying level all means that all messages are printed asynchronously regardless of the severity level. (Optional) For limit number-of-buffers, specify the number of buffers to be queued for the terminal after which new messages are dropped. The range is 0 to 2147483647. The default is 20.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable synchronization of unsolicited messages and debug output, use the **no logging synchronous** [**level severity-level** | **all**] [**limit** *number-of-buffers*] line configuration command.

Enabling and Disabling Time Stamps on Log Messages

By default, log messages are not time stamped.

Beginning in privileged EXEC mode, follow these steps to enable time-stamping of log messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	service timestamps log uptime or service timestamps log datetime [msec] [localtime] [show-timezone]	Enable log time stamps. The first command enables time stamps on log messages, showing the time since the system was rebooted. The second command enables time stamps on log messages. Depending on the options selected, the time stamp can include the date, time in milliseconds relative to the local time zone, and the time-zone name.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable time stamps for both debug and log messages, use the **no service timestamps** global configuration command.

This example shows part of a logging display with the **service timestamps log datetime** global configuration command enabled:

```
*Mar  1 18:46:11: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

This example shows part of a logging display with the **service timestamps log uptime** global configuration command enabled:

```
00:00:46: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
```

Enabling and Disabling Sequence Numbers in Log Messages

Because there is a chance that more than one log message can have the same time stamp, you can display messages with sequence numbers so that you can unambiguously refer to a single message. By default, sequence numbers in log messages are not displayed.

Beginning in privileged EXEC mode, follow these steps to enable sequence numbers in log messages. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	service sequence-numbers	Enable sequence numbers.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable sequence numbers, use the **no service sequence-numbers** global configuration command.

This example shows part of a logging display with sequence numbers enabled:

```
000019: %SYS-5-CONFIG_I: Configured from console by vty2 (10.34.195.36)
```

Defining the Message Severity Level

You can limit messages displayed to the selected device by specifying the severity level of the message, which are described in [Table 24-3](#).

Beginning in privileged EXEC mode, follow these steps to define the message severity level. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging console <i>level</i>	Limit messages logged to the console. By default, the console receives debugging messages and numerically lower levels (see Table 24-3 on page 24-9).
Step 3	logging monitor <i>level</i>	Limit messages logged to the terminal lines. By default, the terminal receives debugging messages and numerically lower levels (see Table 24-3 on page 24-9).
Step 4	logging trap <i>level</i>	Limit messages logged to the syslog servers. By default, syslog servers receive informational messages and numerically lower levels (see Table 24-3 on page 24-9). For complete syslog server configuration steps, see the “ Configuring UNIX Syslog Servers ” section on page 24-10.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config or show logging	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.



Note

Specifying a *level* causes messages at that level and numerically lower levels to appear at the destination.

To disable logging to the console, use the **no logging console** global configuration command. To disable logging to a terminal other than the console, use the **no logging monitor** global configuration command. To disable logging to syslog servers, use the **no logging trap** global configuration command.

[Table 24-3](#) describes the *level* keywords. It also lists the corresponding UNIX syslog definitions from the most severe level to the least severe level.

Table 24-3 Message Logging Level Keywords

Level Keyword	Level	Description	Syslog Definition
emergencies	0	System unstable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

The software generates four other categories of messages:

- Error messages about software or hardware malfunctions, displayed at levels **warnings** through **emergencies**. These types of messages mean that the functionality of the switch is affected. For information on how to recover from these malfunctions, refer to the system message guide for this release.
- Output from the **debug** commands, displayed at the **debugging** level. Debug commands are typically used only by the Technical Assistance Center.
- Interface up or down transitions and system restart messages, displayed at the **notifications** level. This message is only for information; switch functionality is not affected.
- Reload requests and low-process stack messages, displayed at the **informational** level. This message is only for information; switch functionality is not affected.

Limiting Syslog Messages Sent to the History Table and to SNMP

If you enabled syslog message traps to be sent to an SNMP network management station by using the **snmp-server enable trap** global configuration command, you can change the level of messages sent and stored in the switch history table. You also can change the number of messages that are stored in the history table.

Messages are stored in the history table because SNMP traps are not guaranteed to reach their destination. By default, one message of the level **warning** and numerically lower levels (see [Table 24-3 on page 24-9](#)) are stored in the history table even if syslog traps are not enabled.

Beginning in privileged EXEC mode, follow these steps to change the level and history table size defaults. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging history <i>level</i> ¹	Change the default level of syslog messages stored in the history file and sent to the SNMP server. See Table 24-3 on page 24-9 for a list of <i>level</i> keywords. By default, warnings , errors , critical , alerts , and emergencies messages are sent.
Step 3	logging history size <i>number</i>	Specify the number of syslog messages that can be stored in the history table. The default is to store one message. The range is 0 to 500 messages.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

1. [Table 24-3](#) lists the level keywords and severity level. For SNMP usage, the severity level values increase by 1. For example, emergencies equal 1, not 0, and critical equals 3, not 2.

When the history table is full (it contains the maximum number of message entries specified with the **logging history size** global configuration command), the oldest message entry is deleted from the table to allow the new message entry to be stored.

To return the logging of syslog messages to the default level, use the **no logging history** global configuration command. To return the number of messages in the history table to the default value, use the **no logging history size** global configuration command.

Configuring UNIX Syslog Servers

The next sections describe how to configure the UNIX server syslog daemon and how to define the UNIX system logging facility.

Logging Messages to a UNIX Syslog Daemon

Before you can send system log messages to a UNIX syslog server, you must configure the syslog daemon on a UNIX server. This procedure is optional.

Log in as root, and perform these steps:



Note

Some recent versions of UNIX syslog daemons no longer accept by default syslog packets from the network. If this is the case with your system, use the UNIX **man syslogd** command to determine what options must be added to or removed from the syslog command line to enable logging of remote syslog messages.

Step 1 Add a line such as the following to the file `/etc/syslog.conf`:

```
local7.debug /usr/adm/logs/cisco.log
```

The **local7** keyword specifies the logging facility to be used; see [Table 24-4 on page 24-12](#) for information on the facilities. The **debug** keyword specifies the syslog level; see [Table 24-3 on page 24-9](#) for information on the severity levels. The syslog daemon sends messages at this level or at a more severe level to the file specified in the next field. The file must already exist, and the syslog daemon must have permission to write to it.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

```
$ touch /var/log/cisco.log
$ chmod 666 /var/log/cisco.log
```

Step 3 Make sure the syslog daemon reads the new changes:

```
$ kill -HUP `cat /etc/syslog.pid`
```

For more information, see the **man syslog.conf** and **man syslogd** commands on your UNIX system.

Configuring the UNIX System Logging Facility

When sending system log messages to an external device, you can cause the switch to identify its messages as originating from any of the UNIX syslog facilities.

Beginning in privileged EXEC mode, follow these steps to configure UNIX system facility message logging. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging host	Log messages to a UNIX syslog server host by entering its IP address. To build a list of syslog servers that receive logging messages, enter this command more than once.
Step 3	logging trap level	Limit messages logged to the syslog servers. Be default, syslog servers receive informational messages and lower. See Table 24-3 on page 24-9 for <i>level</i> keywords.
Step 4	logging facility facility-type	Configure the syslog facility. See Table 24-4 on page 24-12 for <i>facility-type</i> keywords. The default is local7 .
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a syslog server, use the **no logging host** global configuration command, and specify the syslog server IP address. To disable logging to syslog servers, enter the **no logging trap** global configuration command.

Table 24-4 lists the UNIX system facilities supported by the software. For more information about these facilities, consult the operator's manual for your UNIX operating system.

Table 24-4 Logging Facility-Type Keywords

Facility Type Keyword	Description
auth	Authorization system
cron	Cron facility
daemon	System daemon
kern	Kernel
local0-7	Locally defined messages
lpr	Line printer system
mail	Mail system
news	USENET news
sys9-14	System use
syslog	System log
user	User process
uucp	UNIX-to-UNIX copy system

Displaying the Logging Configuration

To display the logging configuration and the contents of the log buffer, use the **show logging** privileged EXEC command. For information about the fields in this display, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.



Configuring SNMP

This chapter describes how to configure the Simple Network Management Protocol (SNMP) on the Catalyst 2970 switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This chapter consists of these sections:

- [Understanding SNMP, page 25-1](#)
- [Configuring SNMP, page 25-6](#)
- [Displaying SNMP Status, page 25-16](#)

Understanding SNMP

SNMP is an application-layer protocol that provides a message format for communication between managers and agents. The SNMP system consists of an SNMP manager, an SNMP agent, and a management information base (MIB). The SNMP manager can be part of a network management system (NMS) such as CiscoWorks. The agent and MIB reside on the switch. To configure SNMP on the switch, you define the relationship between the manager and the agent.

The SNMP agent contains MIB variables whose values the SNMP manager can request or change. A manager can get a value from an agent or store a value into the agent. The agent gathers data from the MIB, the repository for information about device parameters and network data. The agent can also respond to a manager's requests to get or set data.

An agent can send unsolicited traps to the manager. Traps are messages alerting the SNMP manager to a condition on the network. Traps can mean improper user authentication, restarts, link status (up or down), MAC address tracking, closing of a Transmission Control Protocol (TCP) connection, loss of connection to a neighbor, or other significant events.

This section includes information about these topics:

- [SNMP Versions, page 25-2](#)
- [SNMP Manager Functions, page 25-3](#)
- [SNMP Agent Functions, page 25-4](#)
- [SNMP Community Strings, page 25-4](#)
- [Using SNMP to Access MIB Variables, page 25-5](#)
- [SNMP Notifications, page 25-5](#)
- [SNMP ifIndex MIB Object Values, page 25-6](#)

SNMP Versions

This software release supports these SNMP versions:

- **SNMPv1**—The Simple Network Management Protocol, a Full Internet Standard, defined in RFC 1157.
- **SNMPv2C** replaces the Party-based Administrative and Security Framework of SNMPv2Classic with the community-string-based Administrative Framework of SNMPv2C while retaining the bulk retrieval and improved error handling of SNMPv2Classic. It has these features:
 - **SNMPv2**—Version 2 of the Simple Network Management Protocol, a Draft Internet Standard, defined in RFCs 1902 through 1907.
 - **SNMPv2C**—The community-string-based Administrative Framework for SNMPv2, an Experimental Internet Protocol defined in RFC 1901.
- **SNMPv3**—Version 3 of the SNMP is an interoperable standards-based protocol defined in RFCs 2273 to 2275. SNMPv3 provides secure access to devices by authenticating and encrypting packets over the network and includes these security features:
 - **Message integrity**—ensuring that a packet was not tampered with in transit
 - **Authentication**—determining that the message is from a valid source
 - **Encryption**—mixing the contents of a package to prevent it from being read by an unauthorized source.



Note

To select encryption, enter the **priv** keyword. This keyword is available only when the crypto (encrypted) software image is installed.

Both SNMPv1 and SNMPv2C use a community-based form of security. The community of managers able to access the agent's MIB is defined by an IP address access control list and password.

SNMPv2C includes a bulk retrieval mechanism and more detailed error message reporting to management stations. The bulk retrieval mechanism retrieves tables and large quantities of information, minimizing the number of round-trips required. The SNMPv2C improved error-handling includes expanded error codes that distinguish different kinds of error conditions; these conditions are reported through a single error code in SNMPv1. Error return codes in SNMPv2C report the error type.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy set up for a user and the group within which the user resides. A security level is the permitted level of security within a security model. A combination of the security level and the security model determine which security mechanism is used when handling an SNMP packet. Available security models are SNMPv1, SNMPv2C, and SNMPv3.

Table 25-1 identifies the characteristics of the different combinations of security models and levels.

Table 25-1 SNMP Security Models and Levels

Model	Level	Authentication	Encryption	Result
SNMPv1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv2C	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
SNMPv3	noAuthNoPriv	Username	No	Uses a username match for authentication.
SNMPv3	authNoPriv	MD5 or SHA	No	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms.
SNMPv3	authPriv (requires the cryptographic software image)	MD5 or SHA	DES	Provides authentication based on the HMAC-MD5 or HMAC-SHA algorithms. Provides DES 56-bit encryption in addition to authentication based on the CBC-DES (DES-56) standard.

You must configure the SNMP agent to use the SNMP version supported by the management station. Because an agent can communicate with multiple managers, you can configure the software to support communications using SNMPv1, and SNMPv2C, and SNMPv3 protocols.

SNMP Manager Functions

The SNMP manager uses information in the MIB to perform the operations described in Table 25-2.

Table 25-2 SNMP Operations

Operation	Description
get-request	Retrieves a value from a specific variable.
get-next-request	Retrieves a value from a variable within a table. ¹
get-bulk-request ²	Retrieves large blocks of data, such as multiple rows in a table, that would otherwise require the transmission of many small blocks of data.
get-response	Replies to a get-request, get-next-request, and set-request sent by an NMS.
set-request	Stores a value in a specific variable.
trap	An unsolicited message sent by an SNMP agent to an SNMP manager when some event has occurred.

1. With this operation, an SNMP manager does not need to know the exact variable name. A sequential search is performed to find the needed variable from within a table.
2. The **get-bulk** command only works with SNMPv2 or later.

SNMP Agent Functions

The SNMP agent responds to SNMP manager requests as follows:

- Get a MIB variable—The SNMP agent begins this function in response to a request from the NMS. The agent retrieves the value of the requested MIB variable and responds to the NMS with that value.
- Set a MIB variable—The SNMP agent begins this function in response to a message from the NMS. The SNMP agent changes the value of the MIB variable to the value requested by the NMS.

The SNMP agent also sends unsolicited trap messages to notify an NMS that a significant event has occurred on the agent. Examples of trap conditions include, but are not limited to, when a port or module goes up or down, when spanning-tree topology changes occur, and when authentication failures occur.

SNMP Community Strings

SNMP community strings authenticate access to MIB objects and function as embedded passwords. In order for the NMS to access the switch, the community string definitions on the NMS must match at least one of the three community string definitions on the switch.

A community string can have one of these attributes:

- Read-only (RO)—Gives read access to authorized management stations to all objects in the MIB except the community strings, but does not allow write access
- Read-write (RW)—Gives read and write access to authorized management stations to all objects in the MIB, but does not allow access to the community strings
- Read-write-all—Gives read and write access to authorized management stations to all objects in the MIB, including the community strings



Note

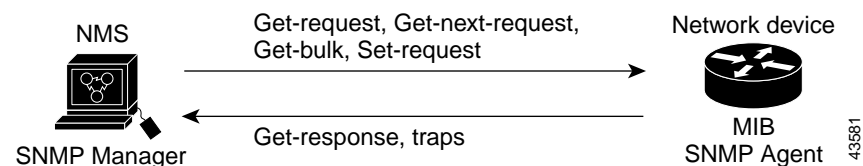
When a cluster is created, the command switch manages the exchange of messages among member switches and the SNMP application. The Cluster Management software appends the member switch number (*@esN*, where *N* is the switch number) to the first configured RW and RO community strings on the command switch and propagates them to the member switches. For more information, see [Chapter 5, “Clustering Switches.”](#)

Using SNMP to Access MIB Variables

An example of an NMS is the CiscoWorks network management software. CiscoWorks 2000 software uses the switch MIB variables to set device variables and to poll devices on the network for specific information. The results of a poll can be displayed as a graph and analyzed to troubleshoot internetworking problems, increase network performance, verify the configuration of devices, monitor traffic loads, and more.

As shown in [Figure 25-1](#), the SNMP agent gathers data from the MIB. The agent can send traps, or notification of certain events, to the SNMP manager, which receives and processes the traps. Traps alert the SNMP manager to a condition on the network such as improper user authentication, restarts, link status (up or down), MAC address tracking, and so forth. The SNMP agent also responds to MIB-related queries sent by the SNMP manager in *get-request*, *get-next-request*, and *set-request* format.

Figure 25-1 *SNMP Network*



For information on supported MIBs and how to access them, see [Appendix A, “Supported MIBs.”](#)

SNMP Notifications

SNMP allows the switch to send notifications to SNMP managers when particular events occur. SNMP notifications can be sent as traps or inform requests. In command syntax, unless there is an option in the command to select either traps or informs, the keyword *traps* refers to either traps or informs, or both. Use the **snmp-server host** command to specify whether to send SNMP notifications as traps or informs.



Note

SNMPv1 does not support informs.

Traps are unreliable because the receiver does not send an acknowledgment when it receives a trap, and the sender cannot determine if the trap was received. When an SNMP manager receives an inform request, it acknowledges the message with an SNMP response protocol data unit (PDU). If the sender does not receive a response, the inform request can be sent again. Because they can be re-sent, informs are more likely than traps to reach their intended destination.

The characteristics that make informs more reliable than traps also consume more resources in the switch and in the network. Unlike a trap, which is discarded as soon as it is sent, an inform request is held in memory until a response is received or the request times out. Traps are sent only once, but an inform might be re-sent or retried several times. The retries increase traffic and contribute to a higher overhead on the network. Therefore, traps and informs require a trade-off between reliability and resources. If it is important that the SNMP manager receive every notification, use inform requests. If traffic on the network or memory in the switch is a concern and notification is not required, use traps.

SNMP ifIndex MIB Object Values

In an NMS, the IF-MIB generates and assigns an interface index (ifIndex) object value that is a unique number greater than zero to identify a physical or a logical interface. When the switch reboots or the switch software is upgraded, the switch uses this same value for the interface. For example, if the switch assigns the Gigabit Ethernet interface 0/5 an ifIndex value of 10003, this value is the same after the switch reboots.

The switch uses one of the values in [Table 25-3](#) to assign an ifIndex value to an interface:

Table 25-3 *ifIndex Values*

Interface Type	ifIndex Range
SVI ¹	1–4999
EtherChannel	5000–5012
Loopback	5013–5077
Tunnel	5078–5142
Physical (such as Gigabit Ethernet or SFP ² -module interfaces)	10000–14500
Null	14501

1. SVI = switch virtual interface
2. SFP = small form-factor pluggable

**Note**

The switch might not use sequential values within a range.

Configuring SNMP

This section describes how to configure SNMP on your switch. It contains this configuration information:

- [Default SNMP Configuration, page 25-7](#)
- [SNMP Configuration Guidelines, page 25-7](#)
- [Disabling the SNMP Agent, page 25-8](#)
- [Configuring Community Strings, page 25-8](#)
- [Configuring SNMP Groups and Users, page 25-9](#)
- [Configuring SNMP Notifications, page 25-11](#)
- [Setting the Agent Contact and Location Information, page 25-14](#)
- [Limiting TFTP Servers Used Through SNMP, page 25-15](#)
- [SNMP Examples, page 25-15](#)

Default SNMP Configuration

Table 25-4 shows the default SNMP configuration.

Table 25-4 Default SNMP Configuration

Feature	Default Setting
SNMP agent	Enabled
SNMP community strings	Read-Only: Public Read-Write: Private Read-Write-all: Secret
SNMP trap receiver	None configured
SNMP traps	None enabled
SNMP version	If no version keyword is present, the default is version 1.
SNMPv3 authentication	If no keyword is entered, the default is the noauth (noAuthNoPriv) security level.
SNMP notification type	If no type is specified, all notifications are sent.

SNMP Configuration Guidelines

An SNMP *group* is a table that maps SNMP users to SNMP views. An SNMP *user* is a member of an SNMP group. An SNMP *host* is the recipient of an SNMP trap operation. An SNMP *engine ID* is a name for the local or remote SNMP engine.

When configuring SNMP, follow these guidelines:

- When configuring an SNMP group, do not specify a notify view. The **snmp-server host** global configuration command autogenerates a notify view for the user and then adds it to the group associated with that user. Modifying the group's notify view affects all users associated with that group. Refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1* for information about when you should configure notify views.
- To configure a remote user, specify the IP address or port number for the remote SNMP agent of the device where the user resides.
- Before you configure remote users for a particular agent, configure the SNMP engine ID, using the **snmp-server engineID** global configuration with the **remote** option. The remote agent's SNMP engine ID and user password are used to compute the authentication and privacy digests. If you do not configure the remote engine ID first, the configuration command fails.
- When configuring SNMP informs, you need to configure the SNMP engine ID for the remote agent in the SNMP database before you can send proxy requests or informs to it.
- Changing the value of the SNMP engine ID has important side effects. A user's password (entered on the command line) is converted to an MD5 or SHA security digest based on the password and the local engine ID. The command-line password is then destroyed, as required by RFC 2274. Because of this deletion, if the value of the engine ID changes, the security digests of SNMPv3 users become invalid, and you need to reconfigure SNMP users by using the **snmp-server user username** global configuration command. Similar restrictions require the reconfiguration of community strings when the engine ID changes.

Disabling the SNMP Agent

Beginning in privileged EXEC mode, follow these steps to disable the SNMP agent:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no snmp-server	Disable the SNMP agent operation.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **no snmp-server** global configuration command disables all running versions (version 1, version 2C, and version 3) on the device. No specific IOS command exists to enable SNMP. The first **snmp-server** global configuration command that you enter enables all versions of SNMP.

Configuring Community Strings

You use the SNMP community string to define the relationship between the SNMP manager and the agent. The community string acts like a password to permit access to the agent on the switch. Optionally, you can specify one or more of these characteristics associated with the string:

- An access list of IP addresses of the SNMP managers that are permitted to use the community string to gain access to the agent
- A MIB view, which defines the subset of all MIB objects accessible to the given community
- Read and write or read-only permission for the MIB objects accessible to the community

Beginning in privileged EXEC mode, follow these steps to configure a community string on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server community <i>string</i> [view <i>view-name</i>] [ro rw] [<i>access-list-number</i>]	Configure the community string. <ul style="list-style-type: none"> • For <i>string</i>, specify a string that acts like a password and permits access to the SNMP protocol. You can configure one or more community strings of any length. • (Optional) For view, specify the view record accessible to the community. • (Optional) Specify either read-only (ro) if you want authorized management stations to retrieve MIB objects, or specify read-write (rw) if you want authorized management stations to retrieve and modify MIB objects. By default, the community string permits read-only access to all objects. • (Optional) For <i>access-list-number</i>, enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.

	Command	Purpose
Step 3	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	<p>(Optional) If you specified an IP standard access list number in Step 2, then create the list, repeating the command as many times as necessary.</p> <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the SNMP managers that are permitted to use the community string to gain access to the agent. (Optional) For <i>source-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. <p>Recall that the access list is always terminated by an implicit deny statement for everything.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

**Note**

To disable access for an SNMP community, set the community string for that community to the null string (do not enter a value for the community string).

To remove a specific community string, use the **no snmp-server community** *string* global configuration command.

This example shows how to assign the string *comaccess* to SNMP, to allow read-only access, and to specify that IP access list 4 can use the community string to gain access to the switch SNMP agent:

```
Switch(config)# snmp-server community comaccess ro 4
```

Configuring SNMP Groups and Users

You can specify an identification name (engine ID) for the local or remote SNMP server engine on the switch. You can configure an SNMP server group that maps SNMP users to SNMP views, and you can add new users to the SNMP group.

Beginning in privileged EXEC mode, follow these steps to configure SNMP on the switch:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server engineID { local <i>engineid-string</i> remote <i>ip-address</i> [udp-port <i>port-number</i>] <i>engineid-string</i> }	<p>Configure a name for either the local or remote copy of SNMP.</p> <ul style="list-style-type: none"> The <i>engineid-string</i> is a 24-character ID string with the name of the copy of SNMP. You need not specify the entire 24-character engine ID if it has trailing zeros. Specify only the portion of the engine ID up to the point where only zeros remain in the value. For example, to configure an engine ID of 123400000000000000000000, you can enter this: snmp-server engineID local 1234 If you select remote, specify the <i>ip-address</i> of the device that contains the remote copy of SNMP and the optional UDP port on the remote device. The default is 162.
Step 3	snmp-server group <i>groupname</i> { v1 v2c v3 [auth noauth priv] } [read <i>readview</i>] [write <i>writeview</i>] [notify <i>notifyview</i>] [access <i>access-list</i>]	<p>Configure a new SNMP group on the remote device.</p> <ul style="list-style-type: none"> For <i>groupname</i>, specify the name of the group. Specify a security model: <ul style="list-style-type: none"> v1 is the least secure of the possible security models. v2c is the second least secure model. It allows transmission of informs and integers twice the normal width. v3, the most secure, requires you to select an authentication level: <ul style="list-style-type: none"> auth—Enables the Message Digest 5 (MD5) and the Secure Hash Algorithm (SHA) packet authentication. noauth—The noAuthNoPriv security level. This is the default if no keyword is specified. priv—Enables Data Encryption Standard (DES) packet encryption (also called <i>privacy</i>). <p>Note The priv keyword is available only when the crypto software image is installed.</p> <ul style="list-style-type: none"> (Optional) Enter read <i>readview</i> with a string (not to exceed 64 characters) that is the name of the view in which you can only view the contents of the agent. (Optional) Enter write <i>writeview</i> with a string (not to exceed 64 characters) that is the name of the view in which you enter data and configure the contents of the agent. (Optional) Enter notify <i>notifyview</i> with a string (not to exceed 64 characters) that is the name of the view in which you specify a notify, inform, or trap. (Optional) Enter access <i>access-list</i> with a string (not to exceed 64 characters) that is the name of the access list.

	Command	Purpose
Step 4	snmp-server user <i>username groupname</i> [remote <i>host</i> [udp-port <i>port</i>]] { v1 v2c v3 [auth { md5 sha } <i>auth-password</i>]} [encrypted] [access <i>access-list</i>]	Configure a new user to an SNMP group. <ul style="list-style-type: none"> The <i>username</i> is the name of the user on the host that connects to the agent. The <i>groupname</i> is the name of the group to which the user is associated. (Optional) Enter remote to specify a remote SNMP entity to which the user belongs and the hostname or IP address of that entity with the optional UDP port number. The default is 162. Enter the SNMP version number (v1, or v2c, or v3). If you enter v3, you have these additional options: <ul style="list-style-type: none"> auth is an authentication level setting session, which can be either the HMAC-MD5-96 or the HMAC-SHA-96 authentication level, and requires a password string (not to exceed 64 characters). encrypted specifies that the password appears in encrypted format. (Optional) Enter access <i>access-list</i> with a string (not to exceed 64 characters) that is the name of the access list.
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Configuring SNMP Notifications

A trap manager is a management station that receives and processes traps. Traps are system alerts that the switch generates when certain events occur. By default, no trap manager is defined, and no traps are sent. Switches running this IOS release can have an unlimited number of trap managers.



Note

Many commands use the word *traps* in the command syntax. Unless there is an option in the command to select either traps or informs, the keyword *traps* refers to either traps, informs, or both. Use the **snmp-server host** command to specify whether to send SNMP notifications as traps or informs.

Table 25-5 describes the supported switch traps (notification types). You can enable any or all of these traps and configure a trap manager to receive them.



Note

Although visible in the command-line interface (CLI) online help, the **fru-ctrl** keyword is not supported.

Table 25-5 Switch Notification Types

Notification Type Keyword	Description
bridge	Generates STP bridge MIB traps.
cluster	Generates a trap when the cluster configuration changes.
config	Generates a trap for SNMP configuration changes.
config-copy	Generates a trap for SNMP copy configuration changes.
entity	Generates a trap for SNMP entity changes.
envmon	Generates environmental monitor traps. You can enable any or all of these environmental traps: fan, shutdown, supply, temperature.
flash	Generates SNMP FLASH notifications.
hsrp	Generates a trap for Hot Standby Router Protocol (HSRP) changes.
mac-notification	Generates a trap for MAC address notifications.
port-security	Generates SNMP port security traps. You can also set a maximum trap rate per second. The range is from 0 to 1000; the default is 0, which means that there is no rate limit.
rtr	Generates a trap for the SNMP Response Time Reporter (RTR).
snmp	Generates a trap for SNMP-type notifications for authentication, cold start, warm start, link up or link down.
stpx	Generates SNMP STP Extended MIB traps.
syslog	Generates SNMP syslog traps.
tty	Generates a trap for TCP connections.
vlan-membership	Generates a trap for SNMP VLAN membership changes.
vlancreate	Generates SNMP VLAN created traps.
vlandelete	Generates SNMP VLAN deleted traps.
vtp	Generates a trap for VLAN Trunking Protocol (VTP) changes.

**Note**

Though visible in the command-line help string, the **fru-ctrl** and flash **insertion** and **removal** keywords are not supported.

You can use the **snmp-server host** global configuration command to a specific host to receive the notification types listed in [Table 25-5](#).

Beginning in privileged EXEC mode, follow these steps to configure the switch to send traps or informs to a host:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server engineID remote <i>ip-address engineid-string</i>	Specify the engine ID for the remote host.
Step 3	snmp-server user <i>username</i> <i>groupname remote host</i> [udp-port <i>port</i>] { v1 v2c v3 [auth { md5 sha } <i>auth-password</i>]} [encrypted] [access <i>access-list</i>]	Configure an SNMP user to be associated with the remote host created in Step 2. Note You cannot configure a remote user for an address without first configuring the engine ID for the remote host. If you try to configure the user before configuring the remote engine ID, you receive an error message, and the command is not executed.
Step 4	snmp-server host <i>host-addr</i> [traps informs] [version { 1 2c 3 [auth noauth priv]}] <i>community-string</i> [udp-port <i>port</i>] [<i>notification-type</i>]	Specify the recipient of an SNMP trap operation. <ul style="list-style-type: none">For <i>host-addr</i>, specify the name or Internet address of the host (the targeted recipient).(Optional) Enter traps (the default) to send SNMP traps to the host.(Optional) Enter informs to send SNMP informs to the host.(Optional) Specify the SNMP version (1, 2c, or 3). SNMPv1 does not support informs.(Optional) For version 3, select authentication level auth, noauth, or priv. Note The priv keyword is available only when the crypto software image is installed. <ul style="list-style-type: none">For <i>community-string</i>, enter the password-like community string sent with the notification operation.(Optional) For udp-port <i>port</i>, enter the remote device UDP port.(Optional) For <i>notification-type</i>, use the keywords listed in Table 25-5 on page 25-12. If no type is specified, all notifications are sent.
Step 5	snmp-server enable traps <i>notification-types</i>	Enable the switch to send traps or informs and specify the type of notifications to be sent. For a list of notification types, see Table 25-5 on page 25-12 , or enter this: snmp-server enable traps ? To enable multiple types of traps, you must enter a separate snmp-server enable traps command for each trap type. Although visible in the CLI online help, the fru-ctrl keyword is not supported.
Step 6	snmp-server trap-source <i>interface-id</i>	(Optional) Specify the source interface, which provides the IP address for the trap message. This command also sets the source IP address for informs.
Step 7	snmp-server queue-length <i>length</i>	(Optional) Establish the message queue length for each trap host. The range is 1 to 1000; the default is 10.

	Command	Purpose
Step 8	snmp-server trap-timeout <i>seconds</i>	(Optional) Define how often to resend trap messages. The range is 1 to 1000; the default is 30 seconds.
Step 9	end	Return to privileged EXEC mode.
Step 10	show running-config	Verify your entries.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

The **snmp-server host** command specifies which hosts receive the notifications. The **snmp-server enable trap** command globally enables the mechanism for the specified notification (for traps and informs). To enable a host to receive an inform, you must configure an **snmp-server host informs** command for the host and globally enable informs by using the **snmp-server enable traps** command.

To remove the specified host from receiving traps, use the **no snmp-server host** *host* global configuration command. The **no snmp-server host** command with no keywords disables traps, but not informs, to the host. To disable informs, use the **no snmp-server host informs** global configuration command. To disable a specific trap type, use the **no snmp-server enable traps** *notification-types* global configuration command.

Setting the Agent Contact and Location Information

Beginning in privileged EXEC mode, follow these steps to set the system contact and location of the SNMP agent so that these descriptions can be accessed through the configuration file:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server contact <i>text</i>	Set the system contact string. For example: snmp-server contact Dial System Operator at beeper 21555.
Step 3	snmp-server location <i>text</i>	Set the system location string. For example: snmp-server location Building 3/Room 222
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Limiting TFTP Servers Used Through SNMP

Beginning in privileged EXEC mode, follow these steps to limit the TFTP servers used for saving and loading configuration files through SNMP to the servers specified in an access list:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server tftp-server-list <i>access-list-number</i>	Limit TFTP servers used for configuration file copies through SNMP to the servers in the access list. For <i>access-list-number</i> , enter an IP standard access list numbered from 1 to 99 and 1300 to 1999.
Step 3	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	Create a standard access list, repeating the command as many times as necessary. <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number specified in Step 2. The deny keyword denies access if the conditions are matched. The permit keyword permits access if the conditions are matched. For <i>source</i>, enter the IP address of the TFTP servers that can access the switch. (Optional) For <i>source-wildcard</i>, enter the wildcard bits, in dotted decimal notation, to be applied to the source. Place ones in the bit positions that you want to ignore. Recall that the access list is always terminated by an implicit deny statement for everything.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

SNMP Examples

This example shows how to enable all versions of SNMP. The configuration permits any SNMP manager to access all objects with read-only permissions using the community string *public*. This configuration does not cause the switch to send any traps.

```
Switch(config)# snmp-server community public
```

This example shows how to permit any SNMP manager to access all objects with read-only permission using the community string *public*. The switch also sends VTP traps to the hosts 192.180.1.111 and 192.180.1.33 using SNMPv1 and to the host 192.180.1.27 using SNMPv2C. The community string *public* is sent with the traps.

```
Switch(config)# snmp-server community public
Switch(config)# snmp-server enable traps vtp
Switch(config)# snmp-server host 192.180.1.27 version 2c public
Switch(config)# snmp-server host 192.180.1.111 version 1 public
Switch(config)# snmp-server host 192.180.1.33 public
```

This example shows how to allow read-only access for all objects to members of access list 4 that use the *comaccess* community string. No other SNMP managers have access to any objects. SNMP Authentication Failure traps are sent by SNMPv2C to the host *cisco.com* using the community string *public*.

```
Switch(config)# snmp-server community comaccess ro 4
Switch(config)# snmp-server enable traps snmp authentication
Switch(config)# snmp-server host cisco.com version 2c public
```

This example shows how to send Entity MIB traps to the host *cisco.com*. The community string is restricted. The first line enables the switch to send Entity MIB traps in addition to any traps previously enabled. The second line specifies the destination of these traps and overwrites any previous **snmp-server host** commands for the host *cisco.com*.

```
Switch(config)# snmp-server enable traps entity
Switch(config)# snmp-server host cisco.com restricted entity
```

This example shows how to enable the switch to send all traps to the host *myhost.cisco.com* using the community string *public*:

```
Switch(config)# snmp-server enable traps
Switch(config)# snmp-server host myhost.cisco.com public
```

Displaying SNMP Status

To display SNMP input and output statistics, including the number of illegal community string entries, errors, and requested variables, use the **show snmp** privileged EXEC command. You can also use the other privileged EXEC commands in [Table 25-6](#) to display SNMP information. For information about the fields in the output displays, refer to the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

Table 25-6 Commands for Displaying SNMP Information

Feature	Default Setting
show snmp	Displays SNMP statistics.
show snmp engineID [local remote]	Displays information on the local SNMP engine and all remote engines that have been configured on the device.
show snmp group	Displays information on each SNMP group on the network.
show snmp user	Displays information on each SNMP user name in the SNMP users table.



Note

Though visible in the command-line help strings, the **fru-ctrl** keyword is not supported. The **snmp-server enable informs** command is not supported. To enable the sending of SNMP inform notifications, use the **snmp-server enable traps** command combined with the **snmp-server host host-addr informs** command.



Configuring Network Security with ACLs

This chapter describes how to configure network security on the Catalyst 2970 switch by using access control lists (ACLs), which are also referred to in commands and tables as access lists.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the “Configuring IP Services” section of the *Cisco IOS IP and IP Routing Configuration Guide* and the *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.

This chapter consists of these sections:

- [Understanding ACLs, page 26-1](#)
- [Configuring IP ACLs, page 26-5](#)
- [Creating Named MAC Extended ACLs, page 26-20](#)
- [Configuring VLAN Maps, page 26-22](#)
- [Displaying ACL Configuration, page 26-29](#)

Understanding ACLs

Packet filtering can help limit network traffic and restrict network use by certain users or devices. ACLs can filter traffic as it passes through a switch and permit or deny packets crossing specified interfaces or VLANs. An ACL is a sequential collection of permit and deny conditions that apply to packets. When a packet is received on an interface, the switch compares the fields in the packet against any applied ACLs to verify that the packet has the required permissions to be forwarded, based on the criteria specified in the access lists. It tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packets. Because the switch stops testing conditions after the first match, the order of conditions in the list is critical. If no conditions match, the switch rejects the packets. If there are no restrictions, the switch forwards the packet; otherwise, the switch drops the packet. The switch can access-control all packets it switches, including packets bridged within a VLAN.

You configure access lists on a switch to provide basic security for your network. If you do not configure ACLs, all packets passing through the switch could be allowed onto all parts of the network. You can use ACLs to control which hosts can access different parts of a network or to decide which types of traffic are forwarded or blocked. For example, you can allow e-mail traffic to be forwarded but not Telnet traffic.

An ACL contains an ordered list of access control entries (ACEs). Each ACE specifies *permit* or *deny* and a set of conditions the packet must satisfy in order to match the ACE. The meaning of *permit* or *deny* depends on the context in which the ACL is used.

The switch supports IP ACLs and Ethernet (MAC) ACLs:

- IP ACLs filter IP traffic, including TCP, User Datagram Protocol (UDP), Internet Group Management Protocol (IGMP), and Internet Control Message Protocol (ICMP).
- Ethernet ACLs filter non-IP traffic.

This switch also supports quality of service (QoS) classification ACLs. For more information, see the [“Classification Based on QoS ACLs” section on page 27-7](#).

This section includes information on these topics:

- [Supported ACLs, page 26-2](#)
- [Handling Fragmented and Unfragmented Traffic, page 26-4](#)

Supported ACLs

The switch supports two applications of ACLs to filter traffic:

- Port ACLs access-control traffic entering a Layer 2 interface. The switch does not support port ACLs in the outbound direction. You can apply only one IP access list and one MAC access list to a Layer 2 interface.
- VLAN ACLs or VLAN maps access-control all packets (bridged and routed). You can use VLAN maps to filter traffic between devices in the same VLAN. VLAN maps are configured to provide access-control based on Layer 3 addresses for IP. Unsupported protocols are access-controlled through MAC addresses using Ethernet ACEs. After a VLAN map is applied to a VLAN, all packets (routed or bridged) entering the VLAN are checked against the VLAN map. Packets can either enter the VLAN through a switch port or through a routed port after being routed.

You can use input port ACLs and VLAN maps on the same switch. However, a port ACL takes precedence over a VLAN map. When an input port ACL is applied to an interface that belongs to a VLAN that has a VLAN map applied, incoming packets received on the interface with the port ACL applied are filtered by the port ACL. Other packets are filtered by the VLAN map.

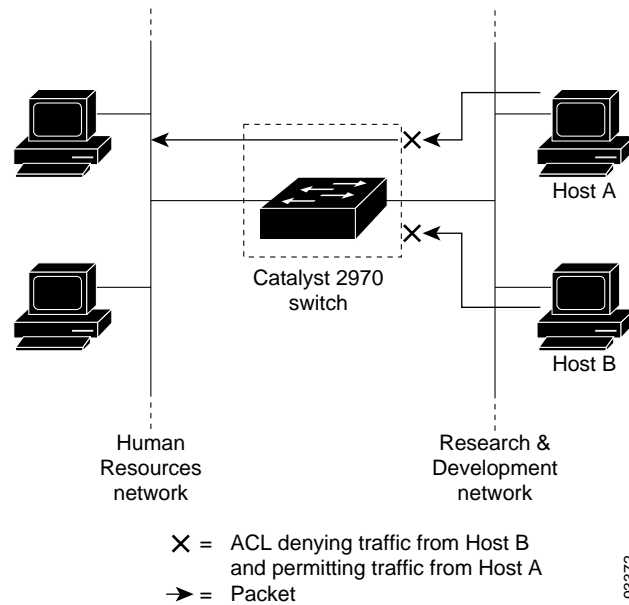
Port ACLs

Port ACLs are ACLs that are applied to Layer 2 interfaces on a switch. Port ACLs are supported only on physical interfaces and not on EtherChannel interfaces. Port ACLs are applied only on interfaces for inbound traffic. These access lists are supported on Layer 2 interfaces:

- Standard IP access lists using source addresses
- Extended IP access lists using source and destination addresses and optional protocol type information
- MAC extended access lists using source and destination MAC addresses and optional protocol type information

The switch examines ACLs associated with features configured on a given interface and permits or denies packet forwarding based on how the packet matches the entries in the ACL. ACLs can only be applied to Layer 2 interfaces in the inbound direction. In the example in [Figure 26-1](#), if all workstations were in the same VLAN, ACLs applied at the Layer 2 input would allow Host A to access the Human Resources network, but prevent Host B from accessing the same network.

Figure 26-1 Using ACLs to Control Traffic to a Network



When you apply a port ACL to a trunk port, the ACL filters traffic on all VLANs present on the trunk port. When you apply a port ACL to a port with voice VLAN, the ACL filters traffic on both data and voice VLANs.

With port ACLs, you can filter IP traffic by using IP access lists and non-IP traffic by using MAC addresses. You can filter both IP and non-IP traffic on the same Layer 2 interface by applying both an IP access list and a MAC access list to the interface.



Note

You cannot apply more than one IP access list and one MAC access list to a Layer 2 interface. If an IP access list or MAC access list is already configured on a Layer 2 interface and you apply a new IP access list or MAC access list to the interface, the new ACL replaces the previously configured one.

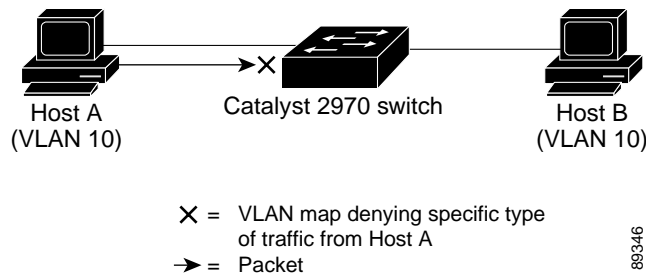
VLAN Maps

You use VLAN ACLs or VLAN maps to filter traffic between devices in the same VLAN. When a VLAN map is applied to a VLAN, all packets being forwarded in the VLAN are checked against the VLAN map. VLAN maps are used for security packet filtering. VLAN maps are not defined by direction (input or output).

You can configure VLAN maps to match Layer 3 addresses for IP traffic. All non-IP protocols are access-controlled through MAC addresses and EtherType using MAC VLAN maps. (IP traffic *is not* access controlled by MAC VLAN maps.) You can enforce VLAN maps only on packets going through the switch; you cannot enforce VLAN maps on traffic between hosts on a hub or on another switch connected to this switch.

With VLAN maps, forwarding of packets is permitted or denied, based on the action specified in the map. Figure 26-2 illustrates how a VLAN map is applied to deny a specific type of traffic from Host A in VLAN 10 from being forwarded. You can apply only one VLAN map to a VLAN.

Figure 26-2 Using VLAN Maps to Control Traffic



Handling Fragmented and Unfragmented Traffic

IP packets can be fragmented as they cross the network. When this happens, only the fragment containing the beginning of the packet contains the Layer 4 information, such as TCP or UDP port numbers, ICMP type and code, and so on. All other fragments are missing this information.

Some ACEs do not check Layer 4 information and therefore can be applied to all packet fragments. ACEs that do test Layer 4 information cannot be applied in the standard manner to most of the fragments in a fragmented IP packet. When the fragment contains no Layer 4 information and the ACE tests some Layer 4 information, the matching rules are modified:

- Permit ACEs that check the Layer 3 information in the fragment (including protocol type, such as TCP, UDP, and so on) are considered to match the fragment regardless of what the missing Layer 4 information might have been.
- Deny ACEs that check Layer 4 information never match a fragment unless the fragment contains Layer 4 information.

Consider access list 102, configured with these commands, applied to three fragmented packets:

```
Switch(config)# access-list 102 permit tcp any host 10.1.1.1 eq smtp
Switch(config)# access-list 102 deny tcp any host 10.1.1.2 eq telnet
Switch(config)# access-list 102 permit tcp any host 10.1.1.2
Switch(config)# access-list 102 deny tcp any any
```



Note

In the first and second ACEs in the examples, the *eq* keyword after the destination address means to test for the TCP-destination-port well-known numbers equaling Simple Mail Transfer Protocol (SMTP) and Telnet, respectively.

- Packet A is a TCP packet from host 10.2.2.2., port 65000, going to host 10.1.1.1 on the SMTP port. If this packet is fragmented, the first fragment matches the first ACE (a permit) as if it were a complete packet because all Layer 4 information is present. The remaining fragments also match the

first ACE, even though they do not contain the SMTP port information, because the first ACE only checks Layer 3 information when applied to fragments. The information in this example is that the packet is TCP and that the destination is 10.1.1.1.

- Packet B is from host 10.2.2.2, port 65001, going to host 10.1.1.2 on the Telnet port. If this packet is fragmented, the first fragment matches the second ACE (a deny) because all Layer 3 and Layer 4 information is present. The remaining fragments in the packet do not match the second ACE because they are missing Layer 4 information. Instead, they match the third ACE (a permit).

Because the first fragment was denied, host 10.1.1.2 cannot reassemble a complete packet, so packet B is effectively denied. However, the later fragments that are permitted will consume bandwidth on the network and resources of host 10.1.1.2 as it tries to reassemble the packet.

- Fragmented packet C is from host 10.2.2.2, port 65001, going to host 10.1.1.3, port ftp. If this packet is fragmented, the first fragment matches the fourth ACE (a deny). All other fragments also match the fourth ACE because that ACE does not check any Layer 4 information and because Layer 3 information in all fragments shows that they are being sent to host 10.1.1.3, and the earlier permit ACEs were checking different hosts.

Configuring IP ACLs

Configuring IP ACLs on the switch is the same as configuring IP ACLs on other Cisco switches and routers. The process is briefly described here. For more detailed information on configuring ACLs, refer to the “Configuring IP Services” chapter in the *Cisco IP and IP Routing Configuration Guide for IOS Release 12.1*. For detailed information about the commands, refer to *Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1*.

The switch does not support these Cisco IOS router ACL-related features:

- Non-IP protocol ACLs (see [Table 26-1 on page 26-6](#)) or bridge-group ACLs
- IP accounting
- Inbound and outbound rate limiting (except with QoS ACLs)
- Reflexive ACLs or dynamic ACLs (except for some specialized dynamic ACLs used by the switch clustering feature)
- ACL logging

These are the steps to use IP ACLs on the switch:

-
- | | |
|---------------|---|
| Step 1 | Create an ACL by specifying an access list number or name and access conditions. |
| Step 2 | Apply the ACL to interfaces or terminal lines. You can also apply standard and extended IP ACLs to VLAN maps. |
-

This section includes the following information:

- [Creating Standard and Extended IP ACLs, page 26-6](#)
- [Applying an IP ACL to a Terminal Line, page 26-16](#)
- [Applying an IP ACL to an Interface, page 26-16](#)
- [Hardware and Software Treatment of IP ACLs, page 26-17](#)
- [IP ACL Configuration Examples, page 26-17](#)

Creating Standard and Extended IP ACLs

This section describes IP ACLs. An ACL is a sequential collection of permit and deny conditions. The switch tests packets against the conditions in an access list one by one. The first match determines whether the switch accepts or rejects the packet. Because the switch stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the switch denies the packet.

The software supports these types of ACLs or access lists for IP:

- Standard IP access lists use source addresses for matching operations.
- Extended IP access lists use source and destination addresses for matching operations and optional protocol-type information for finer granularity of control.

These sections describe access lists and how to create them:

- [Access List Numbers, page 26-6](#)
- [Creating a Numbered Standard ACL, page 26-7](#)
- [Creating a Numbered Extended ACL, page 26-8](#)
- [Creating Named Standard and Extended ACLs, page 26-11](#)
- [Using Time Ranges with ACLs, page 26-13](#)
- [Including Comments in ACLs, page 26-15](#)

Access List Numbers

The number you use to denote your ACL shows the type of access list that you are creating. [Table 26-1](#) lists the access-list number and corresponding access list type and shows whether or not they are supported in the switch. The switch supports IP standard and IP extended access lists, numbers 1 to 199 and 1300 to 2699.

Table 26-1 Access List Numbers

Access List Number	Type	Supported
1–99	IP standard access list	Yes
100–199	IP extended access list	Yes
200–299	Protocol type-code access list	No
300–399	DECnet access list	No
400–499	XNS standard access list	No
500–599	XNS extended access list	No
600–699	AppleTalk access list	No
700–799	48-bit MAC address access list	No
800–899	IPX standard access list	No
900–999	IPX extended access list	No
1000–1099	IPX SAP access list	No
1100–1199	Extended 48-bit MAC address access list	No
1200–1299	IPX summary address access list	No

Table 26-1 Access List Numbers (continued)

Access List Number	Type	Supported
1300–1999	IP standard access list (expanded range)	Yes
2000–2699	IP extended access list (expanded range)	Yes

**Note**

In addition to numbered standard and extended ACLs, you can also create standard and extended named IP ACLs using the supported numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Creating a Numbered Standard ACL

Beginning in privileged EXEC mode, follow these steps to create a numbered standard ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	<p>Define a standard IP access list by using a source address and wildcard.</p> <p>The <i>access-list-number</i> is a decimal number from 1 to 99 or 1300 to 1999.</p> <p>Enter deny or permit to specify whether to deny or permit access if conditions are matched.</p> <p>The <i>source</i> is the source address of the network or host from which the packet is being sent specified as:</p> <ul style="list-style-type: none"> The 32-bit quantity in dotted-decimal format. The keyword any as an abbreviation for <i>source</i> and <i>source-wildcard</i> of 0.0.0.0 255.255.255.255. You do not need to enter a source-wildcard. The keyword host as an abbreviation for source and source-wildcard of <i>source</i> 0.0.0.0. <p>(Optional) The <i>source-wildcard</i> applies wildcard bits to the source.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire ACL. You cannot delete individual ACEs from numbered access lists.

**Note**

When creating an ACL, remember that, by default, the end of the ACL contains an implicit deny statement for all packets that it did not find a match for before reaching the end. With standard access lists, if you omit the mask from an associated IP host address ACL specification, 0.0.0.0 is assumed to be the mask.

This example shows how to create a standard ACL to deny access to IP host 171.69.198.102, permit access to any others, and display the results.

```
Switch (config)# access-list 2 deny host 171.69.198.102
Switch (config)# access-list 2 permit any
Switch(config)# end
Switch# show access-lists
Standard IP access list 2
    deny 171.69.198.102
    permit any
```

The switch always rewrites the order of standard access lists so that entries with **host** matches and entries with matches having a *don't care* mask of 0.0.0.0 are moved to the top of the list, above any entries with non-zero *don't care* masks. Therefore, in **show** command output and in the configuration file, the ACEs do not necessarily appear in the order in which they were entered.

After creating a numbered standard IP ACL, you can apply it to terminal lines (see the [“Applying an IP ACL to a Terminal Line”](#) section on page 26-16), to interfaces (see the [“Applying an IP ACL to an Interface”](#) section on page 26-16), or to VLANs (see the [“Configuring VLAN Maps”](#) section on page 26-22).

Creating a Numbered Extended ACL

Although standard ACLs use only source addresses for matching, you can use extended ACL source and destination addresses for matching operations and optional protocol type information for finer granularity of control. When you are creating ACEs in numbered extended access lists, remember that after you create the ACL, any additions are placed at the end of the list. You cannot reorder the list or selectively add or remove ACEs from a numbered list.

Some protocols also have specific parameters and keywords that apply to that protocol.

These IP protocols are supported (protocol keywords are in parentheses in bold):

Authentication Header Protocol (**ahp**), Enhanced Interior Gateway Routing Protocol (**eigrp**), Encapsulation Security Payload (**esp**), generic routing encapsulation (**gre**), Internet Control Message Protocol (**icmp**), Internet Group Management Protocol (**igmp**), Interior Gateway Routing Protocol (**igrp**), any Interior Protocol (**ip**), IP in IP tunneling (**ipinip**), KA9Q NOS-compatible IP over IP tunneling (**nos**), Open Shortest Path First routing (**ospf**), Payload Compression Protocol (**pcp**), Protocol Independent Multicast (**pim**), Transmission Control Protocol (**tcp**), or User Datagram Protocol (**udp**).

**Note**

ICMP echo-reply cannot be filtered. All other ICMP codes or types can be filtered.

For more details on the specific keywords relative to each protocol, refer to *Cisco IP and IP Routing Command Reference for IOS Release 12.1*.

**Note**

The switch does not support dynamic or reflexive access lists. It also does not support filtering based on the type of service (ToS) minimize-monetary-cost bit.

Supported parameters can be grouped into these categories: TCP, UDP, ICMP, IGMP, or other IP.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2a	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> <i>source source-wildcard</i> <i>destination destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>] Note If you enter a dscp value, you cannot enter tos or precedence . You can enter both a tos and a precedence value with no dscp .	<p>Define an extended IP access list and the access conditions.</p> <p>The <i>access-list-number</i> is a decimal number from 100 to 199 or 2000 to 2699.</p> <p>Enter deny or permit to specify whether to deny or permit the packet if conditions are matched.</p> <p>For <i>protocol</i>, enter the name or number of an IP protocol: ahp, eigrp, esp, gre, icmp, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, tcp, or udp, or an integer in the range 0 to 255 representing an IP protocol number. To match any Internet protocol (including ICMP, TCP, and UDP) use the keyword ip.</p> <p>Note This step includes options for most IP protocols. For additional specific parameters for TCP, UDP, ICMP, and IGMP, see steps 2b through 2e.</p> <p>The <i>source</i> is the number of the network or host from which the packet is sent.</p> <p>The <i>source-wildcard</i> applies wildcard bits to the source.</p> <p>The <i>destination</i> is the network or host number to which the packet is sent.</p> <p>The <i>destination-wildcard</i> applies wildcard bits to the destination.</p> <p>Source, source-wildcard, destination, and destination-wildcard can be specified as:</p> <ul style="list-style-type: none"> • The 32-bit quantity in dotted-decimal format. • The keyword any for 0.0.0.0 255.255.255.255 (any host). • The keyword host for a single host 0.0.0.0. <p>The other keywords are optional and have these meanings:</p> <ul style="list-style-type: none"> • precedence—Enter to match packets with a precedence level specified as a number from 0 to 7 or by name: routine (0), priority (1), immediate (2), flash (3), flash-override (4), critical (5), internet (6), network (7). • fragments—Enter to check non-initial fragments. • tos—Enter to match by type of service level, specified by a number from 0 to 15 or a name: normal (0), max-reliability (2), max-throughput (4), min-delay (8). • time-range—For an explanation of this keyword, see the “Using Time Ranges with ACLs” section on page 26-13. • dscp—Enter to match packets with the DSCP value specified by a number from 0 to 63, or use the question mark (?) to see a list of available values.
or	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> any any [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	<p>In access-list configuration mode, define an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255 and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255.</p> <p>You can use the any keyword in place of source and destination address and wildcard.</p>

	Command	Purpose
or	access-list <i>access-list-number</i> { deny permit } <i>protocol</i> host <i>source</i> host <i>destination</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	Define an extended IP access list using an abbreviation for a source and source wildcard of <i>source</i> 0.0.0.0 and an abbreviation for a destination and destination wildcard of <i>destination</i> 0.0.0.0. You can use the host keyword in place of source and destination wildcard or mask.
Step 2b	access-list <i>access-list-number</i> { deny permit } tcp <i>source</i> <i>source-wildcard</i> [<i>operator</i> <i>port</i>] <i>destination</i> <i>destination-wildcard</i> [<i>operator</i> <i>port</i>] [established] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>] [<i>flag</i>]	(Optional) Define an extended TCP access list and the access conditions. Enter tcp for Transmission Control Protocol. The parameters are the same as those described in Step 2a with these exceptions: (Optional) Enter an <i>operator</i> and <i>port</i> to compare source (if positioned after <i>source source-wildcard</i>) or destination (if positioned after <i>destination destination-wildcard</i>) port. Possible operators include eq (equal), gt (greater than), lt (less than), neq (not equal), and range (inclusive range). Operators require a port number (range requires two port numbers separated by a space). Enter the <i>port</i> number as a decimal number (from 0 to 65535) or the name of a TCP port. To see TCP port names, use the ? or refer to “Configuring IP Services” section of <i>Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1</i> . Use only TCP port numbers or names when filtering TCP. The additional optional keywords have these meanings: <ul style="list-style-type: none"> • established—Enter to match an established connection. This has the same function as matching on the ack or rst flag. • <i>flag</i>—Enter one of these flags to match by the specified TCP header bits: ack (acknowledge), fin (finish), psh (push), rst (reset), syn (synchronize), or urg (urgent).
Step 2c	access-list <i>access-list-number</i> { deny permit } udp <i>source</i> <i>source-wildcard</i> [<i>operator</i> <i>port</i>] <i>destination</i> <i>destination-wildcard</i> [<i>operator</i> <i>port</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended UDP access list and the access conditions. Enter udp for the User Datagram Protocol. The UDP parameters are the same as those described for TCP except that [<i>operator</i> [<i>port</i>]] port number or name must be a UDP port number or name, and the flag and established parameters are not valid for UDP.
Step 2d	access-list <i>access-list-number</i> { deny permit } icmp <i>source</i> <i>source-wildcard</i> <i>destination</i> <i>destination-wildcard</i> [<i>icmp-type</i> / [<i>icmp-type</i> <i>icmp-code</i>] / [<i>icmp-message</i>]] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended ICMP access list and the access conditions. Enter icmp for Internet Control Message Protocol. The ICMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of the ICMP message type and code parameters. These optional keywords have these meanings: <ul style="list-style-type: none"> • <i>icmp-type</i>—Enter to filter by ICMP message type, a number from 0 to 255. • <i>icmp-code</i>—Enter to filter ICMP packets that are filtered by ICMP message type by the ICMP message code, a number from 0 to 255. • <i>icmp-message</i>—Enter to filter ICMP packets by ICMP message type name or ICMP message type and code name. To see a list of ICMP message type names and ICMP message type and code names, use the ? or refer to the “Configuring IP Services” section of <i>Cisco IOS IP and IP Routing Command Reference for IOS Release 12.1</i>.

	Command	Purpose
Step 2e	access-list <i>access-list-number</i> { deny permit } igmp <i>source</i> <i>source-wildcard destination</i> <i>destination-wildcard</i> [<i>igmp-type</i>] [precedence <i>precedence</i>] [tos <i>tos</i>] [fragments] [time-range <i>time-range-name</i>] [dscp <i>dscp</i>]	(Optional) Define an extended IGMP access list and the access conditions. Enter igmp for Internet Group Management Protocol. The IGMP parameters are the same as those described for most IP protocols in Step 2a, with the addition of this optional parameter. <i>igmp-type</i> —To match IGMP message type, enter a number from 0 to 15, or enter the message name (dvmrp , host-query , host-report , pim , or trace).
Step 3	show access-lists [<i>number</i> <i>name</i>]	Verify the access list configuration.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no access-list** *access-list-number* global configuration command to delete the entire access list. You cannot delete individual ACEs from numbered access lists.

This example shows how to create and display an extended access list to deny Telnet access from any host in network 171.69.198.0 to any host in network 172.20.52.0 and permit any others. (The **eq** keyword after the destination address means to test for the TCP destination port number equaling Telnet.)

```
Switch(config)# access-list 102 deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq
telnet
Switch(config)# access-list 102 permit tcp any any
Switch(config)# end
Switch# show access-lists
Extended IP access list 102
    deny tcp 171.69.198.0 0.0.0.255 172.20.52.0 0.0.0.255 eq telnet
    permit tcp any any
```

After an ACL is created, any additions (possibly entered from the terminal) are placed at the end of the list. You cannot selectively add or remove access list entries from a numbered access list.



Note

When you are creating an ACL, remember that, by default, the end of the access list contains an implicit deny statement for all packets if it did not find a match before reaching the end.

After creating a numbered extended ACL, you can apply it to terminal lines (see the [“Applying an IP ACL to a Terminal Line”](#) section on page 26-16), to interfaces (see the [“Applying an IP ACL to an Interface”](#) section on page 26-16), or to VLANs (see the [“Configuring VLAN Maps”](#) section on page 26-22).

Creating Named Standard and Extended ACLs

You can identify IP ACLs with an alphanumeric string (a name) rather than a number. You can use named ACLs to configure more IP access lists in a router than if you were to use numbered access lists. If you identify your access list with a name rather than a number, the mode and command syntax are slightly different. However, not all commands that use IP access lists accept a named access list.



Note

The name you give to a standard or extended ACL can also be a number in the supported range of access list numbers. That is, the name of a standard IP ACL can be 1 to 99; the name of an extended IP ACL can be 100 to 199. The advantage of using named ACLs instead of numbered lists is that you can delete individual entries from a named list.

Consider these guidelines and limitations before configuring named ACLs:

- Not all commands that accept a numbered ACL accept a named ACL. ACLs for packet filters and route filters on interfaces can use a name. VLAN maps also accept a name.
- A standard ACL and an extended ACL cannot have the same name.
- Numbered ACLs are also available, as described in the [“Creating Standard and Extended IP ACLs” section on page 26-6](#).
- You can use standard and extended ACLs (named or numbered) in VLAN maps.

Beginning in privileged EXEC mode, follow these steps to create a standard ACL using names:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip access-list standard <i>name</i>	Define a standard IP access list using a name, and enter access-list configuration mode. Note The name can be a number from 1 to 99.
Step 3	deny { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any } or permit { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any }	In access-list configuration mode, specify one or more conditions denied or permitted to determine if the packet is forwarded or dropped. <ul style="list-style-type: none">• host <i>source</i>—A source and source wildcard of <i>source</i> 0.0.0.0.• any—A source and source wildcard of 0.0.0.0 255.255.255.255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a named standard ACL, use the **no ip access-list standard** *name* global configuration command.

Beginning in privileged EXEC mode, follow these steps to create an extended ACL using names:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip access-list extended <i>name</i>	Define an extended IP access list using a name and enter access-list configuration mode. Note The name can be a number from 100 to 199.
Step 3	{ deny permit } <i>protocol</i> { <i>source</i> [<i>source-wildcard</i>] host <i>source</i> any } { <i>destination</i> [<i>destination-wildcard</i>] host <i>destination</i> any } [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [time-range <i>time-range-name</i>]	In access-list configuration mode, specify the conditions allowed or denied. See the “Creating a Numbered Extended ACL” section on page 26-8 for definitions of protocols and other keywords. <ul style="list-style-type: none">• host <i>source</i>—A source and source wildcard of <i>source</i> 0.0.0.0.• host <i>destination</i>—A destination and destination wildcard of <i>destination</i> 0.0.0.0.• any—A source and source wildcard or destination and destination wildcard of 0.0.0.0 255.255.255.255.

	Command	Purpose
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a named extended ACL, use the **no ip access-list extended** *name* global configuration command.

When you are creating standard extended ACLs, remember that, by default, the end of the ACL contains an implicit deny statement for everything if it did not find a match before reaching the end. For standard ACLs, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

After you create an ACL, any additions are placed at the end of the list. You cannot selectively add ACL entries to a specific ACL. However, you can use **no permit** and **no deny** access-list configuration mode commands to remove entries from a named ACL. This example shows how you can delete individual ACEs from the named access list *border-list*:

```
Switch(config)# ip access-list extended border-list
Switch(config-ext-nacl)# no permit ip host 10.1.1.3 any
```

Being able to selectively remove lines from a named ACL is one reason you might use named ACLs instead of numbered ACLs.

After creating a named ACL, you can apply it to interfaces (see the [“Applying an IP ACL to an Interface” section on page 26-16](#)) or VLANs (see the [“Configuring VLAN Maps” section on page 26-22](#)).

Using Time Ranges with ACLs

You can selectively apply extended ACLs based on the time of day and week by using the **time-range** global configuration command. First, define a time-range name and set the times and the dates or the days of the week in the time range. Then enter the time-range name when applying an ACL to set restrictions to the access list. You can use the time range to define when the permit or deny statements in the ACL are in effect, for example, during a specified time period or on specified days of the week. The **time-range** keyword and argument are referenced in the named and numbered extended ACL task tables in the previous sections, the [“Creating Standard and Extended IP ACLs” section on page 26-6](#), and the [“Creating Named Standard and Extended ACLs” section on page 26-11](#).

Time-based access lists trigger CPU activity because the new configuration of the access list must be merged with other features and the combined configuration loaded into the TCAM. For this reason, you should be careful not to have several access lists configured to take affect in close succession (within a small number of minutes of each other.)



Note

The time range relies on the switch system clock; therefore, you need a reliable clock source. We recommend that you use Network Time Protocol (NTP) to synchronize the switch clock. For more information, see the [“Managing the System Time and Date” section on page 6-1](#).

Beginning in privileged EXEC mode, follow these steps to configure an time-range parameter for an ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	time-range <i>time-range-name</i>	Assign a meaningful name (for example, <i>workhours</i>) to the time range to be created, and enter time-range configuration mode. The name cannot contain a space or quotation mark and must begin with a letter.
Step 3	absolute [start <i>time date</i>] [end <i>time date</i>] or periodic <i>day-of-the-week hh:mm to</i> <i>[day-of-the-week] hh:mm</i> or periodic { weekdays weekend daily } <i>hh:mm to hh:mm</i>	Specify when the function it will be applied to is operational. <ul style="list-style-type: none"> You can use only one absolute statement in the time range. If you configure more than one absolute statement, only the one configured last is executed. You can enter multiple periodic statements. For example, you could configure different hours for weekdays and weekends. Refer to the example configurations.
Step 4	end	Return to privileged EXEC mode.
Step 5	show time-range	Verify the time-range configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Repeat the steps if you have multiple items that you want in effect at different times.

To remove a configured time-range limitation, use the **no time-range** *time-range-name* global configuration command.

This example shows how to configure time ranges for *workhours* and for company holidays and to verify your configuration.

```
Switch(config)# time-range workhours
Switch(config-time-range)# periodic weekdays 8:00 to 12:00
Switch(config-time-range)# periodic weekdays 13:00 to 17:00
Switch(config-time-range)# exit
Switch(config)# time-range new_year_day_2003
Switch(config-time-range)# absolute start 00:00 1 Jan 2003 end 23:59 1 Jan 2003
Switch(config-time-range)# exit
Switch(config)# time-range thanksgiving_2003
Switch(config-time-range)# absolute start 00:00 27 Nov 2003 end 23:59 28 Nov 2003
Switch(config-time-range)# exit
Switch(config)# time-range christmas_2003
Switch(config-time-range)# absolute start 00:00 24 Dec 2003 end 23:59 25 Dec 2003
Switch(config-time-range)# end
Switch# show time-range
time-range entry: christmas_2003 (inactive)
    absolute start 00:00 24 December 2003 end 23:59 25 December 2003
time-range entry: new_year_day_2003 (inactive)
    absolute start 00:00 01 January 2003 end 23:59 01 January 2003
time-range entry: thanksgiving_2003 (inactive)
    absolute start 00:00 27 November 2003 end 23:59 28 November 2003
time-range entry: workhours (inactive)
    periodic weekdays 8:00 to 12:00
    periodic weekdays 13:00 to 17:00
```

To apply a time-range, enter the time-range name in an extended ACL that can implement time ranges. This example shows how to create and verify extended access list 188 that denies TCP traffic from any source to any destination during the defined holiday times and permits all TCP traffic during work hours.

```
Switch(config)# access-list 188 deny tcp any any time-range new_year_day_2003
Switch(config)# access-list 188 deny tcp any any time-range thanksgiving_2003
```

```
Switch(config)# access-list 188 deny tcp any any time-range christmas_2003
Switch(config)# access-list 188 permit tcp any any time-range workhours
Switch(config)# end
Switch# show access-lists
Extended IP access list 188
    deny tcp any any time-range new_year_day_2003 (inactive)
    deny tcp any any time-range thansksgiving_2003 (active)
    deny tcp any any time-range christmas_2003 (inactive)
    permit tcp any any time-range workhours (inactive)
```

This example uses named ACLs to permit and deny the same traffic.

```
Switch(config)# ip access-list extended deny_access
Switch(config-ext-nacl)# deny tcp any any time-range new_year_day_2003
Switch(config-ext-nacl)# deny tcp any any time-range thansksgiving_2003
Switch(config-ext-nacl)# deny tcp any any time-range christmas_2003
Switch(config-ext-nacl)# exit
Switch(config)# ip access-list extended may_access
Switch(config-ext-nacl)# permit tcp any any time-range workhours
Switch(config-ext-nacl)# end
Switch# show ip access-lists
Extended IP access list deny_access
    deny tcp any any time-range new_year_day_2003 (inactive)
    deny tcp any any time-range thansksgiving_2003 (inactive)
    deny tcp any any time-range christmas_2003 (inactive)
Extended IP access list may_access
    permit tcp any any time-range workhours (inactive)
```

Including Comments in ACLs

You can use the **remark** keyword to include comments (remarks) about entries in any IP standard or extended ACL. The remarks make the ACL easier for you to understand and scan. Each remark line is limited to 100 characters.

The remark can go before or after a permit or deny statement. You should be consistent about where you put the remark so that it is clear which remark describes which permit or deny statement. For example, it would be confusing to have some remarks before the associated permit or deny statements and some remarks after the associated statements.

To include a comment for IP numbered standard or extended ACLs, use the **access-list access-list number remark remark** global configuration command. To remove the remark, use the **no** form of this command.

In this example, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

For an entry in a named IP ACL, use the **remark** access-list configuration command. To remove the remark, use the **no** form of this command.

In this example, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp host 171.69.2.88 any eq telnet
```

Applying an IP ACL to a Terminal Line

You can use numbered ACLs to control access to one or more terminal lines. You cannot apply named ACLs to lines. You must set identical restrictions on all the virtual terminal lines because a user can attempt to connect to any of them.

For procedures for applying ACLs to interfaces, see the [“Applying an IP ACL to an Interface” section on page 26-16](#). For applying ACLs to VLANs, see the [“Configuring VLAN Maps” section on page 26-22](#).

Beginning in privileged EXEC mode, follow these steps to restrict incoming and outgoing connections between a virtual terminal line and the addresses in an ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	line [console vty] <i>line-number</i>	Identify a specific line to configure, and enter in-line configuration mode. <ul style="list-style-type: none"> console—Specify the console terminal line. The console port is DCE. vty—Specify a virtual terminal for remote console access. The <i>line-number</i> is the first line number in a contiguous group that you want to configure when the line type is specified. The range is from 0 to 16.
Step 3	access-class <i>access-list-number</i> { in out }	Restrict incoming and outgoing connections between a particular virtual terminal line (into a device) and the addresses in an access list.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Display the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an ACL from a terminal line, use the **no access-class** *access-list-number* {**in** | **out**} line configuration command.

Applying an IP ACL to an Interface

This section describes how to apply IP ACLs to network interfaces. You can apply ACLs only to inbound Layer 2 interfaces. Note these guidelines:

- When controlling access to an interface, you can use a named or numbered ACL.
- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over a VLAN map applied to the VLAN. Incoming packets received on the Layer 2 port are always filtered by the port ACL.

Beginning in privileged EXEC mode, follow these steps to control access to an interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Identify a specific interface for configuration, and enter interface configuration mode.

	Command	Purpose
Step 3	ip access-group { <i>access-list-number</i> / <i>name</i> } { in }	Control access to the specified interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Display the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified access group, use the **no ip access-group** {*access-list-number* | *name*} {**in**} interface configuration command.

This example shows how to apply access list 2 on Gigabit Ethernet interface 0/3 to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet0/3
Router(config-if)# ip access-group 2 in
```

For inbound ACLs, after receiving a packet, the switch checks the packet against the ACL. If the ACL permits the packet, the switch continues to process the packet. If the ACL rejects the packet, the switch discards the packet.

By default, the input interface sends ICMP Unreachable messages whenever a packet is discarded. ICMP Unreachables are normally limited to no more than one every one-half second per input interface, but this can be changed by using the **ip icmp rate-limit unreachable** global configuration command.

When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied to the interface and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Hardware and Software Treatment of IP ACLs

ACL processing is primarily accomplished in hardware, but requires forwarding of some traffic flows to the CPU for software processing. If the hardware reaches its capacity to store ACL configurations, packets are sent to the CPU for forwarding. The forwarding rate for software-forwarded traffic is substantially less than for hardware-forwarded traffic. If ACLs cause large numbers of packets to be sent to the CPU, the switch performance can be negatively affected.

When you enter the **show ip access-lists** privileged EXEC command, the match count displayed does not account for packets that are access controlled in hardware. Use the **show access-lists hardware counters** privileged EXEC command to obtain some basic hardware ACL statistics for switched packets.

IP ACL Configuration Examples

This section provides examples of configuring and applying IP ACLs. For detailed information about compiling ACLs, refer to the *Security Configuration Guide* and the “IP Services” chapter of the *Cisco IOS IP and IP Routing Configuration Guide for IOS Release 12.1*.

This example uses a standard ACL to allow access to the interface to a specific Internet host with the address 172.20.128.64.

```
Switch(config)# access-list 6 permit 172.20.128.64 0.0.0
Switch(config)# end
Switch# show access-lists
Standard IP access list 6
    permit 172.20.128.64 wildcard bits 0.0.0.0
```

```
Switch(config)# interface gigabitethernet 0/3
Switch(config-if)# ip access-group 6 in
```

This example uses an extended ACL to deny traffic coming from port 80 (HTTP). It permits all other types of traffic.

```
Switch(config)# access-list 106 deny tcp any any eq 80
Switch(config)# access-list 106 permit ip any any
Switch(config)# end
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# ip access-group 106 in
```

Numbered ACLs

This ACL accepts addresses on network 36.0.0.0 subnets and denies all packets coming from 56.0.0.0 subnets. The ACL is then applied to packets entering Gigabit Ethernet interface 0/1.

```
Switch(config)# access-list 2 permit 36.0.0.0 0.255.255.255
Switch(config)# access-list 2 deny 56.0.0.0 0.255.255.255
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 2 in
```

Extended ACLs

In this example, suppose that you have a network connected to the Internet, and you want any host on the network to be able to form TCP connections to any host on the Internet. However, you do not want IP hosts to be able to form TCP connections to hosts on your network, except to the mail (SMTP) port of a dedicated mail host.

SMTP uses TCP port 25 on one end of the connection and a random port number on the other end. The same port numbers are used throughout the life of the connection. Mail packets coming in from the Internet have a destination port of 25. Because the secure system of the network always accepts mail connections on port 25, the incoming services are controlled.

```
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 23
Switch(config)# access-list 102 permit tcp any 128.88.0.0 0.0.255.255 eq 25
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group 102 in
```

Named ACLs

The *marketing_group* ACL allows any TCP Telnet traffic to the destination address and wildcard 171.69.0.0 0.0.255.255 and denies any other TCP traffic. It permits any other IP traffic.

```
Switch(config)# ip access-list extended marketing_group
Switch(config-ext-nacl)# permit tcp any 171.69.0.0 0.0.255.255 eq telnet
Switch(config-ext-nacl)# deny tcp any any
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
```

The ACLs are applied to Gigabit Ethernet port 0/4 with the *marketing_group* ACL applied to incoming traffic.

```
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# ip access-group marketing_group in
```

Time Range Applied to an IP ACL

This example denies HTTP traffic on IP on Monday through Friday between the hours of 8:00 a.m. and 6:00 p.m. (18:00). The example allows UDP traffic only on Saturday and Sunday from noon to 8:00 p.m. (20:00).

```
Switch(config)# time-range no-http
Switch(config)# periodic weekdays 8:00 to 18:00
!
Switch(config)# time-range udp-yes
Switch(config)# periodic weekend 12:00 to 20:00
!
Switch(config)# ip access-list extended strict
Switch(config-ext-nacl)# deny tcp any any eq www time-range no-http
Switch(config-ext-nacl)# permit udp any any time-range udp-yes
!
Switch(config-ext-nacl)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip access-group strict in
```

Commented IP ACL Entries

In this example of a numbered ACL, the workstation belonging to Jones is allowed access, and the workstation belonging to Smith is not allowed access:

```
Switch(config)# access-list 1 remark Permit only Jones workstation through
Switch(config)# access-list 1 permit 171.69.2.88
Switch(config)# access-list 1 remark Do not allow Smith workstation through
Switch(config)# access-list 1 deny 171.69.3.13
```

In this example of a numbered ACL, the Winter and Smith workstations are not allowed to browse the web:

```
Switch(config)# access-list 100 remark Do not allow Winter to browse the web
Switch(config)# access-list 100 deny host 171.69.3.85 any eq www
Switch(config)# access-list 100 remark Do not allow Smith to browse the web
Switch(config)# access-list 100 deny host 171.69.3.13 any eq www
```

In this example of a named ACL, the Jones subnet is not allowed access:

```
Switch(config)# ip access-list standard prevention
Switch(config-std-nacl)# remark Do not allow Jones subnet through
Switch(config-std-nacl)# deny 171.69.0.0 0.0.255.255
```

In this example of a named ACL, the Jones subnet is not allowed to use outbound Telnet:

```
Switch(config)# ip access-list extended telnetting
Switch(config-ext-nacl)# remark Do not allow Jones subnet to telnet out
Switch(config-ext-nacl)# deny tcp 171.69.0.0 0.0.255.255 any eq telnet
```

Creating Named MAC Extended ACLs

You can filter non-IP traffic on a VLAN or on a Layer 2 interface by using MAC addresses and named MAC extended ACLs. The procedure is similar to that of configuring other extended named ACLs.

For more information about the supported non-IP protocols in the **mac access-list extended** command, refer to the command reference for this release.



Note

Though visible in the command-line help strings, **appletalk** is not supported as a matching condition for the **deny** and **permit** MAC access-list configuration mode commands.

Beginning in privileged EXEC mode, follow these steps to create a named MAC extended ACL:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac access-list extended <i>name</i>	Define an extended MAC access list using a name.
Step 3	{deny permit} {any host <i>source MAC address</i> / <i>source MAC address mask</i> {any host <i>destination MAC address</i> / <i>destination MAC address mask</i> [type mask lsap <i>lsap mask</i> aarp amber dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat larc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp 0-65535] [cos <i>cos</i>	<p>In extended MAC access-list configuration mode, specify to permit or deny any source MAC address, a source MAC address with a mask, or a specific host source MAC address and any destination MAC address, destination MAC address with a mask, or a specific destination MAC address.</p> <p>(Optional) You can also enter these options:</p> <ul style="list-style-type: none"> type mask—An arbitrary EtherType number of a packet with Ethernet II or SNAP encapsulation in decimal, hex, or octal with optional mask of <i>don't care</i> bits applied to the EtherType before testing for a match. lsap <i>lsap mask</i>—An LSAP number of a packet with 802.2 encapsulation in decimal, hex, or octal with optional mask of <i>don't care</i> bits. aarp amber dec-spanning decnet-iv diagnostic dsm etype-6000 etype-8042 lat larc-sca mop-console mop-dump msdos mumps netbios vines-echo vines-ip xns-idp—A non-IP protocol. cos <i>cos</i>—An IEEE 802.1Q cost of service number from 0 to 7 used to set priority.
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>number</i> <i>name</i>]	Show the access list configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no mac access-list extended** *name* global configuration command to delete the entire ACL. You can also delete individual ACEs from named MAC extended ACLs.

This example shows how to create and display an access list named *mac1*, denying only EtherType DECnet Phase IV traffic, but permitting all other types of traffic.

```
Switch(config)# mac access-list extended mac1
Switch(config-ext-macl)# deny any any decnet-iv
Switch(config-ext-macl)# permit any any
Switch(config-ext-macl)# end
Switch # show access-lists
Extended MAC access list mac1
    deny any any decnet-iv
    permit any any
```

Applying a MAC ACL to a Layer 2 Interface

After you create a MAC ACL, you can apply it to a Layer 2 interface to filter non-IP traffic coming in that interface. When you apply the MAC ACL, consider these guidelines:

- If you apply an ACL to a Layer 2 interface that is a member of a VLAN, the Layer 2 (port) ACL takes precedence over a VLAN map applied to the VLAN. Incoming packets received on the Layer 2 port are always filtered by the port ACL.
- You can apply no more than one IP access list and one MAC access list to the same Layer 2 interface. The IP access list filters only IP packets, and the MAC access list filters non-IP packets.
- A Layer 2 interface can have only one MAC access list. If you apply a MAC access list to a Layer 2 interface that has a MAC ACL configured, the new ACL replaces the previously configured one.

Beginning in privileged EXEC mode, follow these steps to apply a MAC access list to control access to a Layer 2 interface:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Identify a specific interface, and enter interface configuration mode. The interface must be a physical Layer 2 interface (port ACL).
Step 3	mac access-group { <i>name</i> } { in }	Control access to the specified interface by using the MAC access list.
		Note Port ACLs are supported only in the inbound direction.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mac access-group [interface <i>interface-id</i>]	Display the MAC access list applied to the interface or all Layer 2 interfaces.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified access group, use the **no mac access-group** {*name*} interface configuration command.

This example shows how to apply MAC access list *mac1* on Gigabit Ethernet interface 0/3 to filter packets entering the interface:

```
Switch(config)# interface gigabitethernet0/3
Router(config-if)# mac access-group mac1 in
```

**Note**

The **mac access-group** interface configuration command is only valid when applied to a physical Layer 2 interface. You cannot use the command on EtherChannel port channels.

After receiving a packet, the switch checks it against the inbound ACL. If the ACL permits it, the switch continues to process the packet. If the ACL rejects the packet, the switch discards it. When you apply an undefined ACL to an interface, the switch acts as if the ACL has not been applied and permits all packets. Remember this behavior if you use undefined ACLs for network security.

Configuring VLAN Maps

This section describes how to configure VLAN maps, which is the only way to control filtering within a VLAN. VLAN maps have no direction. To filter traffic in a specific direction by using a VLAN map, you need to include an ACL with specific source or destination addresses. If there is a match clause for that type of packet (IP or MAC) in the VLAN map, the default action is to drop the packet if the packet does not match any of the entries within the map. If there is no match clause for that type of packet, the default is to forward the packet.

**Note**

For complete syntax and usage information for the commands used in this section, refer to the command reference for this release.

To create a VLAN map and apply it to one or more VLANs, perform these steps:

- Step 1** Create the standard or extended IP ACLs or named MAC extended ACLs that you want to apply to the VLAN. See the [“Creating Standard and Extended IP ACLs” section on page 26-6](#) and the [“Creating a VLAN Map” section on page 26-23](#).
- Step 2** Enter the **vlan access-map** global configuration command to create a VLAN ACL map entry.
- Step 3** In access map configuration mode, optionally enter an **action—forward** (the default) or **drop**—and enter the **match** command to specify an IP packet or a non-IP packet (with only a known MAC address) and to match the packet against one or more ACLs (standard or extended).

**Note**

If the VLAN map has a match clause for the type of packet (IP or MAC) and the packet does not match the type, the default is to drop the packet. If there is no match clause in the VLAN map for that type of packet, and no action specified, the packet is forwarded.

- Step 4** Use the **vlan filter** global configuration command to apply a VLAN map to one or more VLANs.

This section contains these topics:

- [VLAN Map Configuration Guidelines, page 26-23](#)
- [Creating a VLAN Map, page 26-23](#)
- [Applying a VLAN Map to a VLAN, page 26-26](#)
- [Using VLAN Maps in Your Network, page 26-26](#)

VLAN Map Configuration Guidelines

Follow these guidelines when configuring VLAN maps:

- If there is no ACL configured to deny traffic on an interface and *no* VLAN map is configured, all traffic is permitted.
- Each VLAN map consists of a series of entries. The order of entries in an VLAN map is important. A packet that comes into the switch is tested against the first entry in the VLAN map. If it matches, the action specified for that part of the VLAN map is taken. If there is no match, the packet is tested against the next entry in the map.
- If the VLAN map has at least one match clause for the type of packet (IP or MAC) and the packet does not match any of these match clauses, the default is to drop the packet. If there is no match clause for that type of packet in the VLAN map, the default is to forward the packet.
- The system might take longer to boot if you have configured a very large number of ACLs.
- Logging is not supported for VLAN maps.
- If VLAN map configuration cannot be applied in hardware, all packets in that VLAN must be forwarded by software.
- When a switch has an IP access list or MAC access list applied to a Layer 2 interface, and you apply a VLAN map to a VLAN that the port belongs to, the port ACL takes precedence over the VLAN map.
- See the [“Using VLAN Maps in Your Network” section on page 26-26](#) for configuration examples.

Creating a VLAN Map

Each VLAN map consists of an ordered series of entries. Beginning in privileged EXEC mode, follow these steps to create, add to, or delete a VLAN map entry:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan access-map <i>name</i> [<i>number</i>]	Create a VLAN map, and give it a name and (optionally) a number. The number is the sequence number of the entry within the map. When you create VLAN maps with the same name, numbers are assigned sequentially in increments of 10. When modifying or deleting maps, you can enter the number of the map entry that you want to modify or delete. Entering this command changes to access-map configuration mode.
Step 3	action { drop forward }	(Optional) Set the action for the map entry. The default is to forward.
Step 4	match { ip mac } address { <i>name</i> / <i>number</i> } [<i>name</i> / <i>number</i>]	Match the packet (using either the IP or MAC address) against one or more standard or extended access lists. Note that packets are only matched against access lists of the correct protocol type. IP packets are matched against standard or extended IP access lists. Non-IP packets are only matched against named MAC extended access lists.
Step 5	end	Return to global configuration mode.
Step 6	show running-config	Display the access list configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no vlan access-map** *name* global configuration command to delete a map.

Use the **no vlan access-map** *name number* global configuration command to delete a single sequence entry from within the map.

Use the **no action** access-map configuration command to enforce the default action, which is to forward.

VLAN maps do not use the specific permit or deny keywords. To deny a packet by using VLAN maps, create an ACL that would match the packet, and set the action to drop. A permit in the ACL counts as a match. A deny in the ACL means no match.

Examples of ACLs and VLAN Maps

These examples show how to create ACLs and VLAN maps that for specific purposes.

Example 1

This example shows how to create an ACL and a VLAN map to deny a packet. In the first map, any packets that match the *ip1* ACL (TCP packets) would be dropped. You first create the *ip1* ACL to permit any TCP packet and no other packets. Because there is a match clause for IP packets in the VLAN map, the default action is to drop any IP packet that does not match any of the match clauses.

```
Switch(config)# ip access-list extended ip1
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 10
Switch(config-access-map)# match ip address ip1
Switch(config-access-map)# action drop
```

This example shows how to create a VLAN map to permit a packet. ACL *ip2* permits UDP packets and any packets that match the *ip2* ACL are forwarded. In this map, any IP packets that did not match any of the previous ACLs (that is, packets that are not TCP packets or UDP packets) would get dropped.

```
Switch(config)# ip access-list extended ip2
Switch(config-ext-nacl)# permit udp any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map_1 20
Switch(config-access-map)# match ip address ip2
Switch(config-access-map)# action forward
```

Example 2

In this example, the VLAN map has a default action of drop for IP packets and a default action of forward for MAC packets. Used with standard ACL 101 and extended named access lists **igmp-match** and **tcp-match**, the map will have the following results:

- Forward all UDP packets
- Drop all IGMP packets
- Forward all TCP packets
- Drop all other IP packets
- Forward all non-IP packets

```
Switch(config)# access-list 101 permit udp any any
Switch(config)# ip access-list extended igmp-match
Switch(config-ext-nacl)# permit igmp any any
Switch(config)# ip access-list extended tcp-match
Switch(config-ext-nacl)# permit tcp any any
Switch(config-ext-nacl)# exit
```



```
Switch(config)# vlan access-map drop-ip-default 10
Switch(config-access-map)# match ip address 101
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 20
Switch(config-access-map)# match ip address igmp-match
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-ip-default 30
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
```

Example 3

In this example, the VLAN map has a default action of drop for MAC packets and a default action of forward for IP packets. Used with MAC extended access lists **good-hosts** and **good-protocols**, the map will have the following results:

- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Forward MAC packets with decnet-iv or vines-ip protocols
- Drop all other non-IP packets
- Forward all IP packets

```
Switch(config)# mac access-list extended good-hosts
Switch(config-ext-macl)# permit host 000.0c00.0111 any
Switch(config-ext-macl)# permit host 000.0c00.0211 any
Switch(config-ext-nacl)# exit
Switch(config)# mac access-list extended good-protocols
Switch(config-ext-macl)# permit any any decnet-ip
Switch(config-ext-macl)# permit any any vines-ip
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map drop-mac-default 10
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-mac-default 20
Switch(config-access-map)# match mac address good-protocols
Switch(config-access-map)# action forward
```

Example 4

In this example, the VLAN map has a default action of drop for all packets (IP and non-IP). Used with access lists **tcp-match** and **good-hosts** from Examples 2 and 3, the map will have the following results:

- Forward all TCP packets
- Forward MAC packets from hosts 0000.0c00.0111 and 0000.0c00.0211
- Drop all other IP packets
- Drop all other MAC packets

```
Switch(config)# vlan access-map drop-all-default 10
Switch(config-access-map)# match ip address tcp-match
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
Switch(config)# vlan access-map drop-all-default 20
Switch(config-access-map)# match mac address good-hosts
Switch(config-access-map)# action forward
```

Applying a VLAN Map to a VLAN

Beginning in privileged EXEC mode, follow these steps to apply a VLAN map to one or more VLANs:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	vlan filter <i>mapname</i> vlan-list <i>list</i>	Apply the VLAN map to one or more VLAN IDs. The list can be a single VLAN ID (22), a consecutive list (10-22), or a string of VLAN IDs (12, 22, 30). Spaces around the comma and hyphen are optional.
Step 3	show running-config	Display the access list configuration.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the VLAN map, use the **no vlan filter** *mapname* **vlan-list** *list* global configuration command.

This example shows how to apply VLAN map 1 to VLANs 20 through 22:

```
Switch(config)# vlan filter map 1 vlan-list 20-22
```

Using VLAN Maps in Your Network

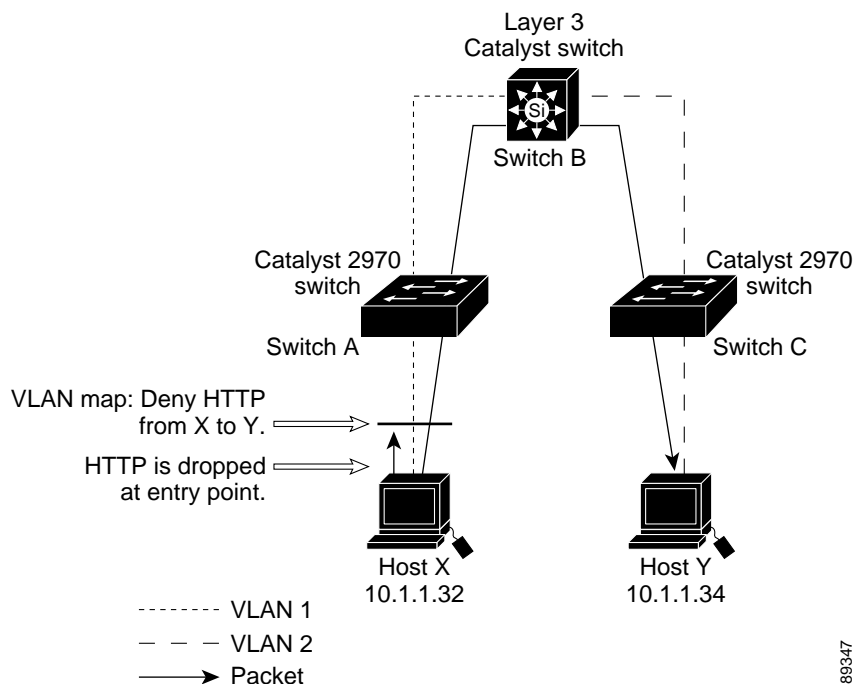
This section describes some typical uses for VLAN maps and includes these topics:

- [Wiring Closet Configuration, page 26-26](#)
- [Denying Access to a Server on a VLAN, page 26-28](#)

Wiring Closet Configuration

In a wiring closet configuration, the switch can support a VLAN map and a QoS classification ACL. In [Figure 26-3](#), assume that Host X and Host Y are in different VLANs and are connected to wiring closet switches A and C. Traffic from Host X to Host Y is eventually being routed by Switch B, which has routing enabled. Traffic from Host X to Host Y can be access-controlled at the traffic entry point, Switch A.

Figure 26-3 Wiring Closet Configuration



89347

If you do not want HTTP traffic switched from Host X to Host Y, you can configure a VLAN map on Switch A to drop all HTTP traffic from Host X (IP address 10.1.1.32) to Host Y (IP address 10.1.1.34) at Switch A and not bridge it to Switch B.

First, define the IP access list *http* that permits (matches) any TCP traffic on the HTTP port.

```
Switch(config)# ip access-list extended http
Switch(config-ext-nacl)# permit tcp host 10.1.1.32 host 10.1.1.34 eq www
Switch(config-ext-nacl)# exit
```

Next, create VLAN access map *map2* so that traffic that matches the *http* access list is dropped and all other IP traffic is forwarded.

```
Switch(config)# vlan access-map map2 10
Switch(config-access-map)# match ip address http
Switch(config-access-map)# action drop
Switch(config-access-map)# exit
Switch(config)# ip access-list extended match_all
Switch(config-ext-nacl)# permit ip any any
Switch(config-ext-nacl)# exit
Switch(config)# vlan access-map map2 20
Switch(config-access-map)# match ip address match_all
Switch(config-access-map)# action forward
```

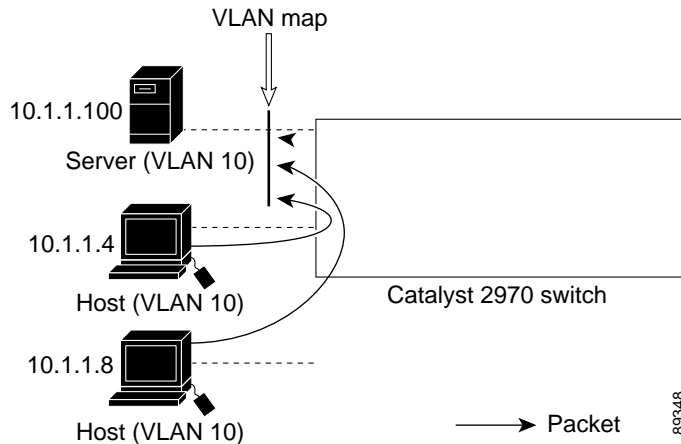
Then, apply VLAN access map *map2* to VLAN 1.

```
Switch(config)# vlan filter map2 vlan 1
```

Denying Access to a Server on a VLAN

You can restrict access to a server on a VLAN. For example, server 10.1.1.100 in VLAN 10 needs to have access denied to hosts 10.1.1.4 and 10.1.1.8 (see [Figure 26-4](#)).

Figure 26-4 Deny Access to a Server on Another VLAN



This example shows how to deny access to a server on another VLAN by creating the VLAN map SERVER1 that denies access to hosts in subnet 10.1.2.0.8, host 10.1.1.4, and host 10.1.1.8 and permits other IP traffic. The final step is to apply the map SERVER1 to VLAN 10.

- Step 1** Define the IP ACL that will match the correct packets.

```
Switch(config)# ip access-list extended SERVER1_ACL
Switch(config-ext-nacl)# permit ip host 10.1.1.4 host 10.1.1.100
Switch(config-ext-nacl)# permit ip host 10.1.1.8 host 10.1.1.100
Switch(config-ext-nacl)# exit
```

- Step 2** Define a VLAN map using this ACL that will drop IP packets that match SERVER1_ACL and forward IP packets that do not match the ACL.

```
Switch(config)# vlan access-map SERVER1_MAP
Switch(config-access-map)# match ip address SERVER1_ACL
Switch(config-access-map)# action drop
Switch(config)# vlan access-map SERVER1_MAP 20
Switch(config-access-map)# action forward
Switch(config-access-map)# exit
```

- Step 3** Apply the VLAN map to VLAN 10.

```
Switch(config)# vlan filter SERVER1_MAP vlan-list 10.
```

Displaying ACL Configuration

You can display the ACLs that are configured on the switch, and you can display the ACLs that have been applied to interfaces and VLANs.

When you use the **ip access-group** interface configuration command to apply ACLs to a Layer 2 interface, you can display the access groups on the interface. You can also display the MAC ACLs applied to a Layer 2 interface. You can use the privileged EXEC commands as described in [Table 26-2](#) to display this information.

Table 26-2 Commands for Displaying Access Lists and Access Groups

Command	Purpose
show access-lists [<i>number / name</i>]	Display the contents of one or all current IP and MAC address access lists or a specific access list (numbered or named).
show ip access-lists [<i>number / name</i>]	Display the contents of all current IP access lists or a specific IP access list (numbered or named).
show running-config [interface <i>interface-id</i>]	Displays the contents of the configuration file for the switch or the specified interface, including all configured MAC and IP access lists and which access groups are applied to an interface.
show mac access-group [interface <i>interface-id</i>]	Displays MAC access lists applied to all Layer 2 interfaces or the specified Layer 2 interface.

You can also display information about VLAN access maps or VLAN filters. Use the privileged EXEC commands in [Table 26-3](#) to display VLAN map information.

Table 26-3 Commands for Displaying VLAN Map Information

Command	Purpose
show vlan access-map [<i>mapname</i>]	Show information about all VLAN access-maps or the specified access map.
show vlan filter [access-map <i>name</i> / vlan <i>vlan-id</i>]	Show information about all VLAN filters or about a specified VLAN or VLAN access map.



Configuring QoS

This chapter describes how to configure quality of service (QoS) by using automatic QoS (auto-QoS) commands or by using standard QoS commands on the Catalyst 2970 switch. With QoS, you can provide preferential treatment to certain traffic at the expense of others. Without QoS, the switch offers best-effort service to each packet, regardless of the packet contents or size. It sends the packets without any assurance of reliability, delay bounds, or throughput.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference this release.

This chapter consists of these sections:

- [Understanding QoS, page 27-1](#)
- [Configuring Auto-QoS, page 27-18](#)
- [Displaying Auto-QoS Information, page 27-26](#)
- [Configuring Standard QoS, page 27-26](#)
- [Displaying Standard QoS Information, page 27-65](#)

Understanding QoS

Typically, networks operate on a best-effort delivery basis, which means that all traffic has equal priority and an equal chance of being delivered in a timely manner. When congestion occurs, all traffic has an equal chance of being dropped.

When you configure the QoS feature, you can select specific network traffic, prioritize it according to its relative importance, and use congestion-management and congestion-avoidance techniques to provide preferential treatment. Implementing QoS in your network makes network performance more predictable and bandwidth utilization more effective.

The QoS implementation is based on the Differentiated Services (Diff-Serv) architecture, an emerging standard from the Internet Engineering Task Force (IETF). This architecture specifies that each packet is classified upon entry into the network.

The classification is carried in the IP packet header, using 6 bits from the deprecated IP type of service (TOS) field to carry the classification (*class*) information. Classification can also be carried in the Layer 2 frame. These special bits in the Layer 2 frame or a Layer 3 packet are described here and shown in Figure 27-1:

- Prioritization bits in Layer 2 frames:

Layer 2 Inter-Switch Link (ISL) frame headers have a 1-byte User field that carries an IEEE 802.1P class of service (CoS) value in the three least-significant bits. On interfaces configured as Layer 2 ISL trunks, all traffic is in ISL frames.

Layer 2 802.1Q frame headers have a 2-byte Tag Control Information field that carries the CoS value in the three most-significant bits, which are called the User Priority bits. On interfaces configured as Layer 2 802.1Q trunks, all traffic is in 802.1Q frames except for traffic in the native VLAN.

Other frame types cannot carry Layer 2 CoS values.

Layer 2 CoS values range from 0 for low priority to 7 for high priority.

- Prioritization bits in Layer 3 packets:

Layer 3 IP packets can carry either an IP precedence value or a Differentiated Services Code Point (DSCP) value. QoS supports the use of either value because DSCP values are backward-compatible with IP precedence values.

IP precedence values range from 0 to 7.

DSCP values range from 0 to 63.

Figure 27-1 QoS Classification Layers in Frames and Packets

Encapsulated Packet

Layer 2 header	IP header	Data
----------------	-----------	------

Layer 2 ISL Frame

ISL header (26 bytes)	Encapsulated frame 1... (24.5 KB)	FCS (4 bytes)
--------------------------	--------------------------------------	------------------

↑ 3 bits used for CoS

Layer 2 802.1Q/P Frame

Preamble	Start frame delimiter	DA	SA	Tag	PT	Data	FCS
----------	--------------------------	----	----	-----	----	------	-----

↑ 3 bits used for CoS (user priority)

Layer 3 IPv4 Packet

Version length	ToS (1 byte)	Len	ID	Offset	TTL	Proto	FCS	IP-SA	IP-DA	Data
-------------------	-----------------	-----	----	--------	-----	-------	-----	-------	-------	------

↑ IP precedence or DSCP

46974



Note

Layer 3 IPv6 packets are treated as non-IP packets and are bridged by the switch.

All switches and routers that access the Internet rely on the class information to provide the same forwarding treatment to packets with the same class information and different treatment to packets with different class information. The class information in the packet can be assigned by end hosts or by switches or routers along the way, based on a configured policy, detailed examination of the packet, or both. Detailed examination of the packet is expected to happen closer to the edge of the network so that the core switches and routers are not overloaded with this task.

Switches and routers along the path can use the class information to limit the amount of resources allocated per traffic class. The behavior of an individual device when handling traffic in the DiffServ architecture is called per-hop behavior. If all devices along a path provide a consistent per-hop behavior, you can construct an end-to-end QoS solution.

Implementing QoS in your network can be a simple or complex task and depends on the QoS features offered by your internetworking devices, the traffic types and patterns in your network, and the granularity of control that you need over incoming and outgoing traffic.

Basic QoS Model

To implement QoS, the switch must distinguish packets or flow from one another (classify), assign a label to indicate the given quality of service as the packets move through the switch, make the packets comply with the configured resource usage limits (police and mark), and provide different treatment (queue and schedule) in all situations where resource contention exists. The switch also needs to ensure that traffic sent from it meets a specific traffic profile (shape).

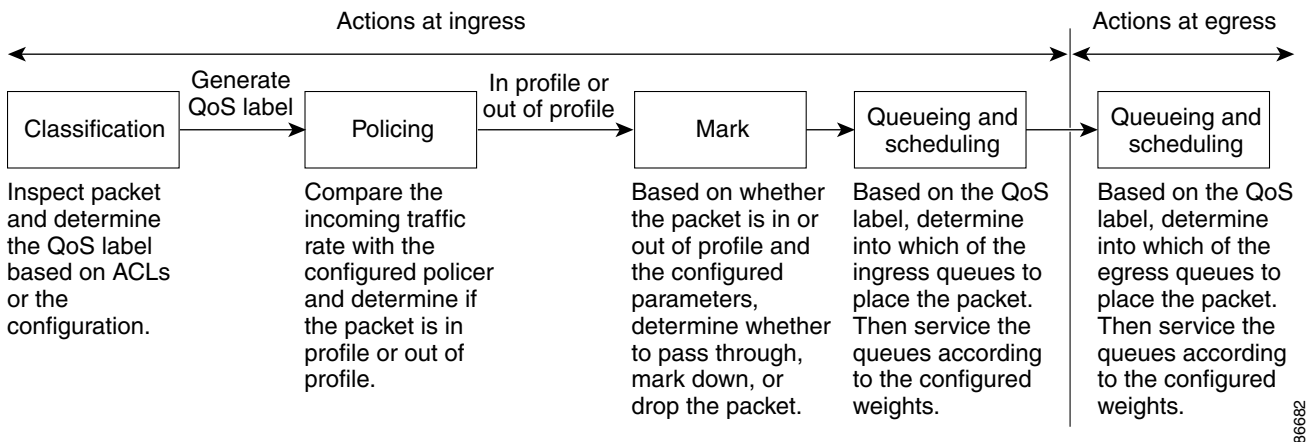
Figure 27-2 shows the basic QoS model. Actions at the ingress interface include classifying traffic, policing, marking, queueing, and scheduling:

- Classification is the process of generating a distinct path for a packet by associating it with a QoS label. The switch maps the CoS or DSCP in the packet to a QoS label to distinguish one kind of traffic from another. The QoS label that is generated identifies all future QoS actions to be performed on this packet. For more information, see the [“Classification” section on page 27-4](#).
- Policing determines whether a packet is in or out of profile by comparing the rate of the incoming traffic to the configured policer. The policer limits the bandwidth consumed by a flow of traffic. The result of this determination is passed to the marker. For more information, see the [“Policing and Marking” section on page 27-8](#).
- Marking evaluates the policer and configuration information for the action to be taken when a packet is out of profile and decides what to do with the packet (pass through a packet without modification, mark down the QoS label in the packet, or drop the packet). For more information, see the [“Policing and Marking” section on page 27-8](#).
- Queueing evaluates the QoS label and the corresponding DSCP or CoS value to determine into which of the two ingress queues to place a packet. Queueing is enhanced with the weighted tail-drop (WTD) algorithm, a congestion-avoidance mechanism. If the threshold is exceeded, the packet is dropped. For more information, see the [“Queueing and Scheduling Overview” section on page 27-11](#).
- Scheduling services the queues based on their configured shaped round robin (SRR) weights. One of the ingress queues is the priority queue, and SRR services it for its configured share before servicing the other queue. For more information, see the [“SRR Shaping and Sharing” section on page 27-12](#).

Actions at the egress interface include queueing and scheduling:

- Queueing evaluates the QoS label and the corresponding DSCP or CoS value to determine into which of the four egress queues to place a packet. Because congestion can occur when multiple ingress ports simultaneously send data to an egress port, WTD is used to differentiate traffic classes and to subject the packets to different thresholds based on the QoS label. If the threshold is exceeded, the packet is dropped. For more information, see the [“Queueing and Scheduling Overview” section on page 27-11](#).
- Scheduling services the four egress queues based on their configured SRR shared or shaped weights. One of the queues (queue 1) can be the expedite queue, which is serviced until empty before the other queues are serviced.

Figure 27-2 Basic QoS Model



Classification

Classification is the process of distinguishing one kind of traffic from another by examining the fields in the packet. Classification is enabled only if QoS is globally enabled on the switch. By default, QoS is globally disabled, so no classification occurs.



Note

Classification occurs only on a physical interface basis. No support exists for classifying packets at the VLAN or the switch virtual interface level.

During classification, the switch performs a lookup and assigns a QoS label to the packet. The QoS label identifies all QoS actions to be performed on the packet and from which queue the packet is sent.

The QoS label is based on the DSCP or the CoS value in the packet and determines the queueing and scheduling actions to perform on the packet. The label is mapped according to the trust setting and the packet type as shown in [Figure 27-3 on page 27-6](#).

You specify which fields in the frame or packet that you want to use to classify incoming traffic. For non-IP traffic, you have these classification options as shown in [Figure 27-3](#):

- Trust the CoS value in the incoming frame (configure the port to trust CoS). Then use the configurable CoS-to-DSCP map to generate a DSCP value for the packet. Layer 2 ISL frame headers carry the CoS value in the three least-significant bits of the 1-byte User field. Layer 2 802.1Q frame headers carry the CoS value in the three most-significant bits of the Tag Control Information field. CoS values range from 0 for low priority to 7 for high priority.
- Trust the DSCP or trust IP precedence value in the incoming frame. These configurations are meaningless for non-IP traffic. If you configure a port with either of these options and non-IP traffic is received, the switch assigns a CoS value and generates a DSCP value from the CoS-to-DSCP map.
- Perform the classification based on a configured Layer 2 MAC access control list (ACL), which can examine the MAC source address, the MAC destination address, and other fields. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or CoS value to assign to the incoming frame.

For IP traffic, you have these classification options as shown in [Figure 27-3](#):

- Trust the DSCP value in the incoming packet (configure the port to trust DSCP), and assign the same DSCP value to the packet. The IETF defines the six most-significant bits of the 1-byte TOS field as the DSCP. The priority represented by a particular DSCP value is configurable. DSCP values range from 0 to 63.

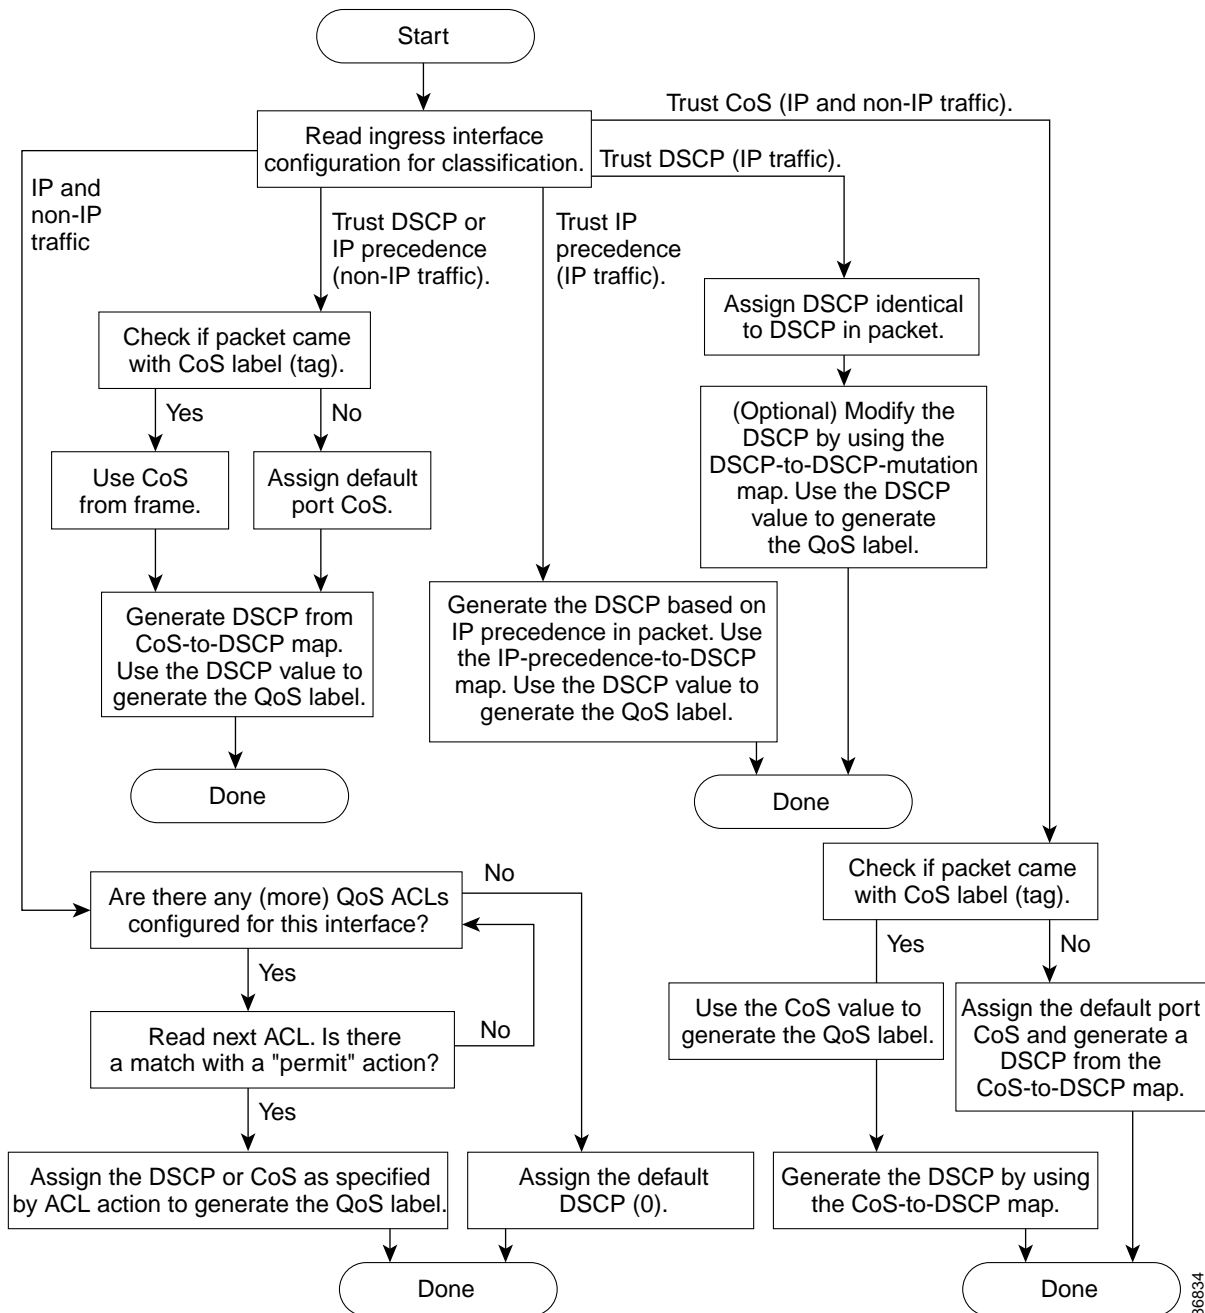
For ports that are on the boundary between two QoS administrative domains, you can modify the DSCP to another value by using the configurable DSCP-to-DSCP-mutation map.

- Trust the IP precedence value in the incoming packet (configure the port to trust IP precedence), and generate a DSCP value for the packet by using the configurable IP-precedence-to-DSCP map. The IP version 4 specification defines the three most-significant bits of the 1-byte ToS field as the IP precedence. IP precedence values range from 0 for low priority to 7 for high priority.
- Trust the CoS value (if present) in the incoming packet, and generate a DSCP value for the packet by using the CoS-to-DSCP map. If the CoS value is not present, use the default port CoS value.
- Perform the classification based on a configured IP standard or an extended ACL, which examines various fields in the IP header. If no ACL is configured, the packet is assigned 0 as the DSCP and CoS values, which means best-effort traffic. Otherwise, the policy-map action specifies a DSCP or CoS value to assign to the incoming frame.

For information on the maps described in this section, see the [“Mapping Tables” section on page 27-10](#). For configuration information on port trust states, see the [“Configuring Classification Using Port Trust States” section on page 27-30](#).

After classification, the packet is sent to the policing, marking, and the ingress queueing and scheduling stages.

Figure 27-3 Classification Flowchart



86834

Classification Based on QoS ACLs

You can use IP standard, IP extended, or Layer 2 MAC ACLs to define a group of packets with the same characteristics (*class*). In the QoS context, the permit and deny actions in the access control entries (ACEs) have different meanings than with security ACLs:

- If a match with a permit action is encountered (first-match principle), the specified QoS-related action is taken.
- If a match with a deny action is encountered, the ACL being processed is skipped, and the next ACL is processed.
- If no match with a permit action is encountered and all the ACEs have been examined, no QoS processing occurs on the packet, and the switch offers best-effort service to the packet.
- If multiple ACLs are configured on an interface, the lookup stops after the packet matches the first ACL with a permit action, and QoS processing begins.



Note

When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.

After a traffic class has been defined with the ACL, you can attach a policy to it. A policy might contain multiple classes with actions specified for each one of them. A policy might include commands to classify the class as a particular aggregate (for example, assign a DSCP) or rate-limit the class. This policy is then attached to a particular port on which it becomes effective.

You implement IP ACLs to classify IP traffic by using the **access-list** global configuration command; you implement Layer 2 MAC ACLs to classify non-IP traffic by using the **mac access-list extended** global configuration command. For configuration information, see the [“Configuring a QoS Policy” section on page 27-36](#).

Classification Based on Class Maps and Policy Maps

A class map is a mechanism that you use to name a specific traffic flow (or class) and to isolate it from all other traffic. The class map defines the criteria used to match against a specific traffic flow to further classify it. The criteria can include matching the access group defined by the ACL or matching a specific list of DSCP or IP precedence values. If you have more than one type of traffic that you want to classify, you can create another class map and use a different name. After a packet is matched against the class-map criteria, you further classify it through the use of a policy map.

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; or specifying the traffic bandwidth limitations and the action to take when the traffic is out of profile. Before a policy map can be effective, you must attach it to an interface.

You create a class map by using the **class-map** global configuration command or the **class** policy-map configuration command. You should use the **class-map** command when the map is shared among many ports. When you enter the **class-map** command, the switch enters the class-map configuration mode. In this mode, you define the match criterion for the traffic by using the **match** class-map configuration command.

You create and name a policy map by using the **policy-map** global configuration command. When you enter this command, the switch enters the policy-map configuration mode. In this mode, you specify the actions to take on a specific traffic class by using the **class**, **trust**, or **set** policy-map configuration and policy-map class configuration commands.

The policy map can contain the **police** and **police aggregate** policy-map class configuration commands, which define the policer, the bandwidth limitations of the traffic, and the action to take if the limits are exceeded.

To make the policy map effective, you attach it to an interface by using the **service-policy** interface configuration command.

For more information, see the “Policing and Marking” section on page 27-8. For configuration information, see the “Configuring a QoS Policy” section on page 27-36.

Policing and Marking

After a packet is classified and has a DSCP-based or CoS-based QoS label assigned to it, the policing and marking process can begin as shown in [Figure 27-4](#).

Policing involves creating a policer that specifies the bandwidth limits for the traffic. Packets that exceed the limits are *out of profile* or *nonconforming*. Each policer determines on a packet-by-packet basis whether the packet is in or out of profile and specifies the actions on the packet. These actions, carried out by the marker, include passing through the packet without modification, dropping the packet, or modifying (marking down) the assigned DSCP of the packet and allowing the packet to pass through. The configurable policed-DSCP map provides the packet with a new DSCP-based QoS label. For information on the policed-DSCP map, see the “Mapping Tables” section on page 27-10. Marked-down packets use the same queues as the original QoS label to prevent packets in a flow from getting out of order.



Note

All traffic, regardless of whether it is bridged or routed, is subjected to a policer, if one is configured. As a result, bridged packets might be dropped or might have their DSCP or CoS fields modified when they are policed and marked.

You can create these types of policers:

- Individual

QoS applies the bandwidth limits specified in the policer separately to each matched traffic class. You configure this type of policer within a policy map by using the **police** policy-map class configuration command.

- Aggregate

QoS applies the bandwidth limits specified in an aggregate policer cumulatively to all matched traffic flows. You configure this type of policer by specifying the aggregate policer name within a policy map by using the **police aggregate** policy-map class configuration command. You specify the bandwidth limits of the policer by using the **mls qos aggregate-policer** global configuration command. In this way, the aggregate policer is shared by multiple classes of traffic within a policy map.

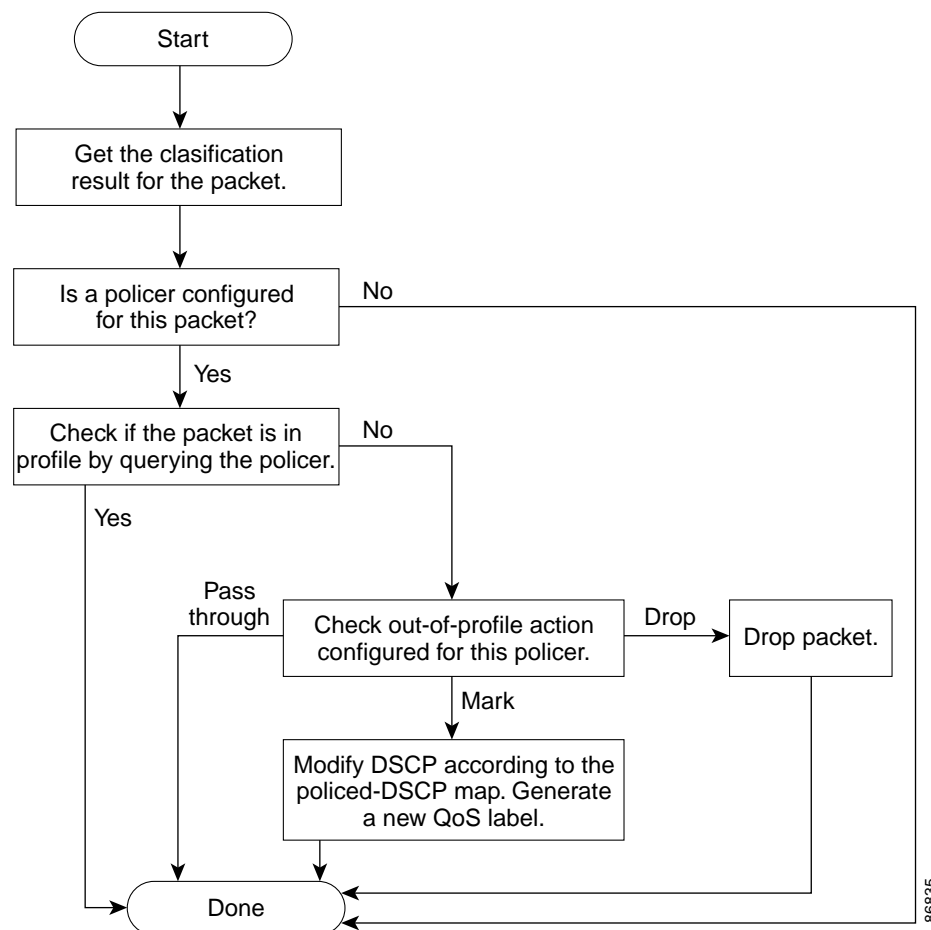
Policing uses a token-bucket algorithm. As each frame is received by the switch, a token is added to the bucket. The bucket has a hole in it and leaks at a rate that you specify as the average traffic rate in bits per second. Each time a token is added to the bucket, the switch performs a check to determine if there is enough room in the bucket. If there is not enough room, the packet is marked as nonconforming, and the specified policer action is taken (dropped or marked down).

How quickly the bucket fills is a function of the bucket depth (burst-byte), the rate at which the tokens are removed (rate-bps), and the duration of the burst above the average rate. The size of the bucket imposes an upper limit on the burst length and determines the number of frames that can be transmitted back-to-back. If the burst is short, the bucket does not overflow, and no action is taken against the traffic flow. However, if a burst is long and at a higher rate, the bucket overflows, and the policing actions are taken against the frames in that burst.

You configure the bucket depth (the maximum burst that is tolerated before the bucket overflows) by using the *burst-byte* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command. You configure how fast (the average rate) that the tokens are removed from the bucket by using the *rate-bps* option of the **police** policy-map class configuration command or the **mls qos aggregate-policer** global configuration command.

After you configure the policy map and policing actions, attach the policy to an ingress interface by using the **service-policy** interface configuration command. For configuration information, see the [“Classifying, Policing, and Marking Traffic by Using Policy Maps”](#) section on page 27-42 and the [“Classifying, Policing, and Marking Traffic by Using Aggregate Policers”](#) section on page 27-45.

Figure 27-4 Policing and Marking Flowchart



Mapping Tables

During QoS processing, the switch represents the priority of all traffic (including non-IP traffic) with an QoS label based on the DSCP or CoS value from the classification stage:

- During classification, QoS uses configurable mapping tables to derive a corresponding DSCP or CoS value from a received CoS, DSCP, or IP precedence value. These maps include the CoS-to-DSCP map and the IP-precedence-to-DSCP map. You configure these maps by using the **mls qos map cos-dscp** and the **mls qos map ip-prec-dscp** global configuration commands.

On an ingress interface configured in the DSCP-trusted state, if the DSCP values are different between the QoS domains, you can apply the configurable DSCP-to-DSCP-mutation map to the interface that is on the boundary between the two QoS domains. You configure this map by using the **mls qos map dscp-mutation** global configuration command.

- During policing, QoS can assign another DSCP value to an IP or a non-IP packet (if the packet is out of profile and the policer specifies a marked-down value). This configurable map is called the policed-DSCP map. You configure this map by using the **mls qos map policed-dscp** global configuration command.
- Before the traffic reaches the scheduling stage, QoS stores the packet in an ingress and an egress queue according to the QoS label. The QoS label is based on the DSCP or the CoS value in the packet and selects the queue through the DSCP input and output queue threshold maps or through the CoS input and output queue threshold maps. You configure these maps by using the **mls qos srr-queue {input | output} dscp-map** and the **mls qos srr-queue {input | output} cos-map** global configuration commands.

The CoS-to-DSCP, DSCP-to-CoS, and the IP-precedence-to-DSCP maps have default values that might or might not be appropriate for your network.

The default DSCP-to-DSCP-mutation map and the default policed-DSCP map are null maps; they map an incoming DSCP value to the same DSCP value. The DSCP-to-DSCP-mutation map is the only map you apply to a specific port. All other maps apply to the entire switch.

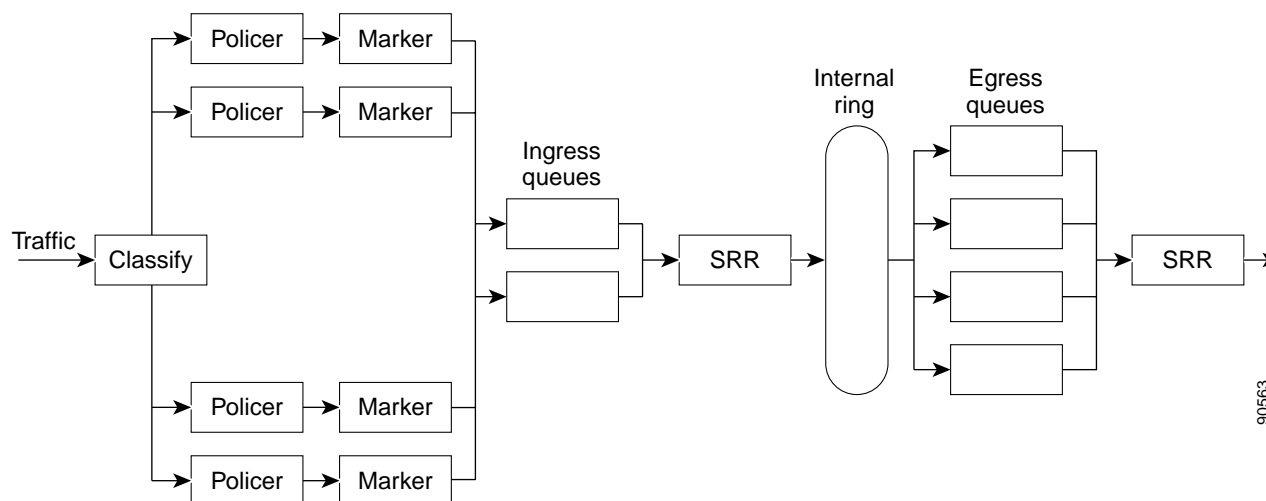
For configuration information, see the [“Configuring DSCP Maps” section on page 27-47](#).

For information about the DSCP and CoS input queue threshold maps, see the [“Queueing and Scheduling on Ingress Queues” section on page 27-13](#). For information about the DSCP and CoS output queue threshold maps, see the [“Queueing and Scheduling on Egress Queues” section on page 27-15](#).

Queueing and Scheduling Overview

The switch has queues at specific points to help prevent congestion as shown in [Figure 27-5](#).

Figure 27-5 Ingress and Egress Queue Location



Because the total ingress bandwidth of all ports can exceed the bandwidth of the internal ring, ingress queues are located after the packet is classified, policed, and marked and before packets are forwarded into the switch fabric. Because multiple ingress ports can simultaneously send packets to an egress port and cause congestion, egress queues are located after the internal ring.

Weighted Tail Drop

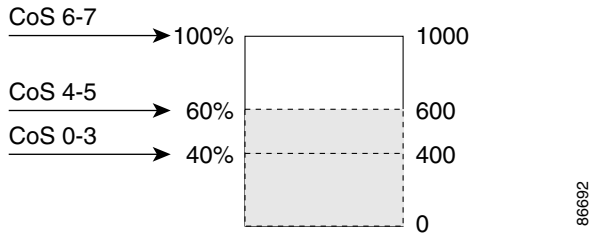
Both the ingress and egress queues use an enhanced version of the tail-drop congestion-avoidance mechanism called weighted tail drop (WTD). WTD is implemented on queues to manage the queue lengths and to provide drop precedences for different traffic classifications.

As a frame is enqueued to a particular queue, WTD uses the frame's assigned QoS label to subject it to different thresholds. If the threshold is exceeded for that QoS label (the space available in the destination queue is less than the size of the frame), the switch drops the frame.

[Figure 27-6](#) shows an example of WTD operating on a queue whose size is 1000 frames. Three drop percentages are configured: 40 percent (400 frames), 60 percent (600 frames), and 100 percent (1000 frames). These percentages mean that up to 400 frames can be queued at the 40-percent threshold, up to 600 frames at the 60-percent threshold, and up to 1000 frames at the 100-percent threshold.

In this example, CoS values 6 and 7 have a greater importance than the other CoS values, and they are assigned to the 100-percent drop threshold (queue-full state). CoS values 4 and 5 are assigned to the 60-percent threshold, and CoS values 0 to 3 are assigned to the 40-percent threshold.

Suppose the queue is already filled with 600 frames, and a new frame arrives. It contains CoS values 4 and 5 and is subjected to the 60-percent threshold. If this frame is added to the queue, the threshold will be exceeded, so the switch drops it.

Figure 27-6 WTD and Queue Operation

For more information, see the [“Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds”](#) section on page 27-53, the [“Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set”](#) section on page 27-58, and the [“Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID”](#) section on page 27-60.

SRR Shaping and Sharing

Both the ingress and egress queues are serviced by SRR, which determines the rate at which packets are sent. On the ingress queues, SRR sends packets to the internal ring. On the egress queues, SRR sends packets to the egress interface.

You can configure SRR on egress queues for sharing or for shaping. However, for ingress queues, sharing is the default mode, and it is the only mode supported.

In shaped mode, the egress queues are guaranteed a percentage of the bandwidth, and they are rate-limited to that amount. Shaped traffic does not use more than the allocated bandwidth even if the link is idle. Shaping provides a more even flow of traffic over time and reduces the peaks and valleys of bursty traffic. With shaping, the absolute value of each weight is used to compute the bandwidth available for the queues.

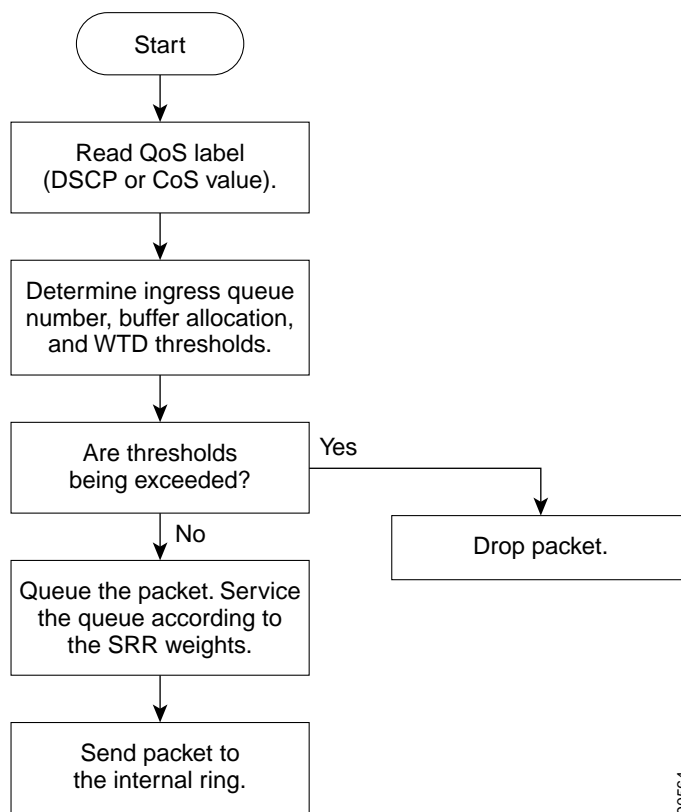
In shared mode, the queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue is empty and no longer requires a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights determines the frequency of dequeuing; the absolute values are meaningless.

For more information, see the [“Allocating Bandwidth Between the Ingress Queues”](#) section on page 27-55, the [“Configuring SRR Shaped Weights on Egress Queues”](#) section on page 27-61, and the [“Configuring SRR Shared Weights on Egress Queues”](#) section on page 27-62.

Queueing and Scheduling on Ingress Queues

Figure 27-7 shows the queueing and scheduling flowchart for ingress ports.

Figure 27-7 Queueing and Scheduling Flowchart for Ingress Ports



Note

SRR services the priority queue for its configured share before servicing the other queue.

The switch supports two configurable ingress queues, which are serviced by SRR in shared mode only. Table 27-1 describes the queues.

Table 27-1 Ingress Queue Types

Queue Type ¹	Function
Normal	User traffic that is considered to be normal priority. You can configure three different thresholds to differentiate among the flows. You can use the mls qos srr-queue input threshold , the mls qos srr-queue input dscp-map , and the mls qos srr-queue input cos-map global configuration commands.
Expedite	High-priority user traffic such as differentiated services (DF) expedited forwarding or voice traffic. You can configure the bandwidth required for this traffic as a percentage of the total traffic by using the mls qos srr-queue input priority-queue global configuration command. The expedite queue has guaranteed bandwidth.

1. The switch uses two nonconfigurable queues for traffic that is essential for proper network operation.

You assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an ingress queue and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue input dscp-map queue *queue-id* {*dscp1...dscp8* | **threshold** *threshold-id dscp1...dscp8*}** or the **mls qos srr-queue input cos-map queue *queue-id* {*cos1...cos8* | **threshold** *threshold-id cos1...cos8*}** global configuration command. You can display the DSCP input queue threshold map and the CoS input queue threshold map by using the **show mls qos maps** privileged EXEC command.

WTD Thresholds

The queues use WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two explicit WTD threshold percentages for threshold ID 1 and ID 2 to the ingress queues by using the **mls qos srr-queue input threshold *queue-id threshold-percentage1 threshold-percentage2*** global configuration command. Each threshold value is a percentage of the total number of allocated buffers for the queue. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot modify it. For more information about how WTD works, see the [“Weighted Tail Drop” section on page 27-11](#).

Buffer and Bandwidth Allocation

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues by using the **mls qos srr-queue input buffers *percentage1 percentage2*** global configuration command. The buffer allocation together with the bandwidth allocation determine how much data can be buffered and sent before packets are dropped. You allocate bandwidth as a percentage by using the **mls qos srr-queue input bandwidth *weight1 weight2*** global configuration command. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue.

Priority Queueing

You can configure one ingress queue as the priority queue by using the **mls qos srr-queue input **priority-queue** *queue-id* **bandwidth** *weight*** global configuration command. The priority queue should be used for traffic (such as voice) that requires guaranteed delivery because this queue is guaranteed part of the bandwidth regardless of the load on the internal ring.

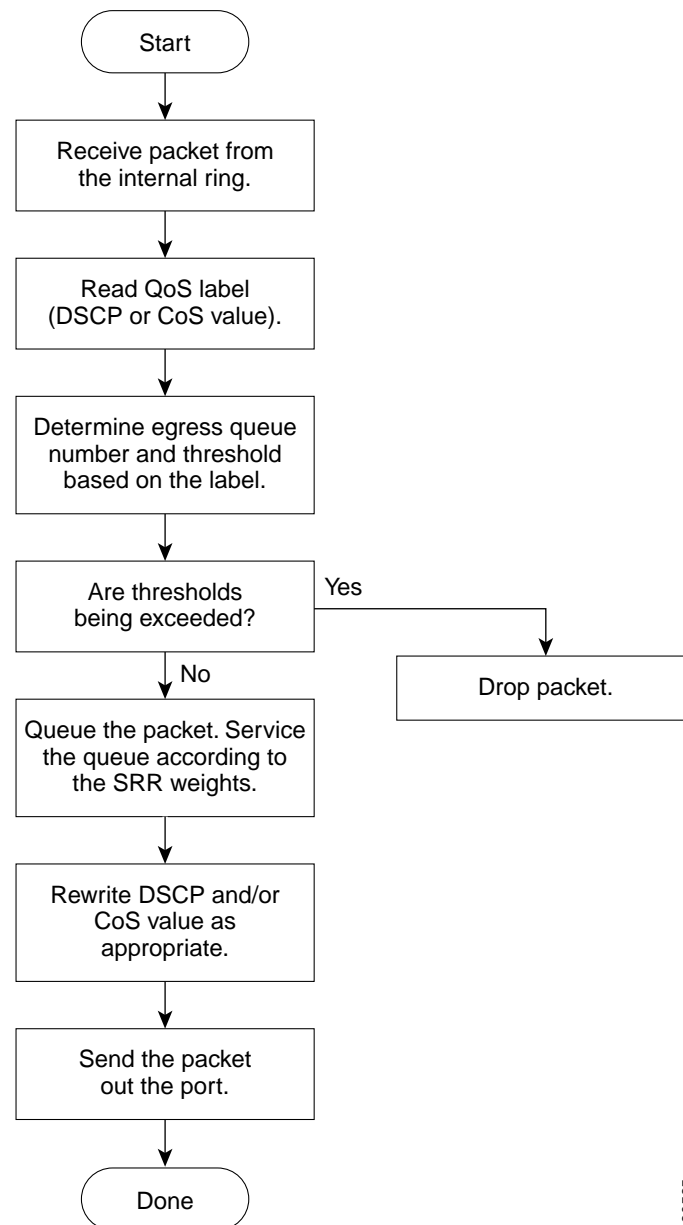
SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input **priority-queue** *queue-id* **bandwidth** *weight*** global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth *weight1 weight2*** global configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the [“Configuring Ingress Queue Characteristics” section on page 27-52](#).

Queueing and Scheduling on Egress Queues

Figure 27-8 shows the queueing and scheduling flowchart for egress ports.

Figure 27-8 Queueing and Scheduling Flowchart for Egress Ports



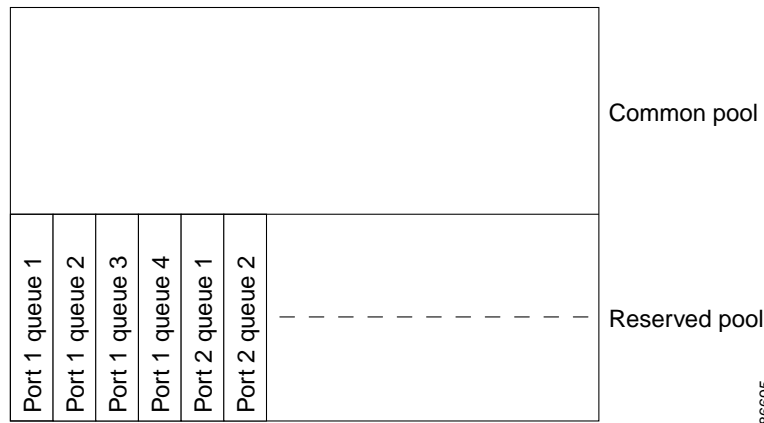
Note

If the expedite queue is enabled, SRR services it until it is empty before servicing the other three queues.

Each port supports four egress queues, one of which (queue 1) can be the egress expedite queue. These queues are assigned to a queue-set. All traffic exiting the switch flows through one of these four queues and is subjected to a threshold based on the QoS label assigned to the packet.

Figure 27-9 shows the egress queue buffer. The buffer space is divided between the common pool and the reserved pool. The switch uses a buffer allocation scheme to reserve a minimum amount of buffers for each egress queue, to prevent any queue or port from consuming all the buffers and depriving other queues, and to determine whether to grant buffer space to a requesting queue. The switch determines whether the target queue has not consumed more buffers than its reserved amount (under-limit), whether it has consumed all of its maximum buffers (over limit), and whether the common pool is empty (no free buffers) or not empty (free buffers). If the queue is not over-limit, the switch can allocate buffer space from the reserved pool or from the common pool (if it is not empty). If there are no free buffers in the common pool or if the queue is over-limit, the switch drops the frame.

Figure 27-9 Egress Queue Buffer Allocation



Buffer and Memory Allocation

You guarantee the availability of buffers, set drop thresholds, and configure the maximum memory allocation for a queue-set by using the **mls qos queue-set output *qset-id* threshold *queue-id* drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold** global configuration command. Each threshold value is a percentage of the queue's allocated memory, which you specify by using the **mls qos queue-set output *qset-id* buffers allocation1 ... allocation4** global configuration command. The sum of all the allocated buffers represents the reserved pool, and the remaining buffers are part of the common pool.

Through buffer allocation, you can ensure that high-priority traffic is buffered. For example, if the buffer space is 400, you can allocate 70 percent of it to queue 1 and 10 percent to queues 2 through 4. Queue 1 then has 280 buffers allocated to it, and queues 2 through 4 each have 40 buffers allocated to them.

You can guarantee that the allocated buffers are reserved for a specific queue in a queue-set. For example, if there are 100 buffers for a queue, you can reserve 50 percent (50 buffers). The switch returns the remaining 50 buffers to the common pool. You also can enable a queue in the full condition to obtain more buffers than are reserved for it by setting a maximum threshold. The switch can allocate the needed buffers from the common pool if the common pool is not empty.

WTD Thresholds

You can assign each packet that flows through the switch to a queue and to a threshold. Specifically, you map DSCP or CoS values to an egress queue and map DSCP or CoS values to a threshold ID. You use the **mls qos srr-queue output dscp-map queue *queue-id* {*dscp1...dscp8* | threshold *threshold-id* *dscp1...dscp8*}** or the **mls qos srr-queue output cos-map queue *queue-id* {*cos1...cos8* | threshold**

`threshold-id cos1...cos8` global configuration command. You can display the DSCP output queue threshold map and the CoS output queue threshold map by using the **show mls qos maps** privileged EXEC command.

The queues use WTD to support distinct drop percentages for different traffic classes. Each queue has three drop thresholds: two configurable (*explicit*) WTD thresholds and one nonconfigurable (*implicit*) threshold preset to the queue-full state. You assign the two WTD threshold percentages for threshold ID 1 and ID 2. The drop threshold for threshold ID 3 is preset to the queue-full state, and you cannot modify it. For more information about how WTD works, see the [“Weighted Tail Drop” section on page 27-11](#).

Shaped or Shared Mode

SRR services each queue-set in shared or shaped mode. You map an interface to a queue-set by using the **queue-set** *qset-id* interface configuration command. You assign shared or shaped weights to the interface by using the **srr-queue bandwidth share** *weight1 weight2 weight3 weight4* or the **srr-queue bandwidth shape** *weight1 weight2 weight3 weight4* interface configuration command. For an explanation of the differences between shaping and sharing, see the [“SRR Shaping and Sharing” section on page 27-12](#).

The buffer allocation together with the SRR weight ratios determine how much data can be buffered and sent before packets are dropped. The weight ratio is the ratio of the frequency in which the SRR scheduler sends packets from each queue.

All four queues participate in the SRR unless the expedite queue is enabled, in which case the first bandwidth weight is ignored and is not used in the ratio calculation. The expedite queue is a priority queue, and it is serviced until empty before the other queues are serviced. You enable the expedite queue by using the **priority-queue out** interface configuration command.

You can combine the commands described in this section to prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues, by allocating a large queue size or by servicing the queue more frequently, and by adjusting queue thresholds so that packets with lower priorities are dropped. For configuration information, see the [“Configuring Egress Queue Characteristics” section on page 27-57](#).



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Packet Modification

A packet is classified, policed, and queued to provide QoS. Packet modifications can occur during this process:

- For IP and non-IP packets, classification involves assigning a QoS label to a packet based on the DSCP or CoS of the received packet. However, the packet is not modified at this stage; only an indication of the assigned DSCP or CoS value is carried along. The reason for this is that QoS classification and forwarding lookups occur in parallel, and it is possible that the packet is forwarded with its original DSCP to the CPU where it is again processed through software.
- During policing, IP and non-IP packets can have another DSCP assigned to them (if they are out of profile and the policer specifies a markdown DSCP). Once again, the DSCP in the packet is not modified, but an indication of the marked-down value is carried along. For IP packets, the packet modification occurs at a later stage; for non-IP packets the DSCP is converted to CoS and used for queueing and scheduling decisions.

- Depending on the QoS label assigned to a frame and the mutation chosen, the DSCP and CoS values of the frame are rewritten. If you do not configure the mutation map and if you configure the interface to trust the DSCP of the incoming frame, the DSCP value in the frame is not changed, but the CoS is rewritten according to the DSCP-to-CoS map. If you configure the interface to trust the CoS of the incoming frame and it is an IP packet, the CoS value in the frame is not changed, but the DSCP might be changed according to the CoS-to-DSCP map.

The input mutation causes the DSCP to be rewritten depending on the new value of DSCP chosen. The set action in a policy map also causes the DSCP to be rewritten.

Configuring Auto-QoS

You can use the auto-QoS feature to simplify the deployment of existing QoS features. Auto-QoS makes assumptions about the network design, and as a result, the switch can prioritize different traffic flows and appropriately use the ingress and egress queues instead of using the default QoS behavior. (The default is that QoS is disabled. The switch then offers best-effort service to each packet, regardless of the packet contents or size, and sends it from a single queue.)

When you enable auto-QoS, it automatically classifies traffic based on the traffic type and ingress packet label. The switch uses the resulting classification to choose the appropriate egress queue.

You use auto-QoS commands to identify ports connected to Cisco IP Phones and to identify ports that receive trusted voice over IP (VoIP) traffic through an uplink. Auto-QoS then performs these functions:

- Detects the presence or absence of Cisco IP Phones
- Configures QoS classification
- Configures egress queues

These sections describe how to configure auto-QoS on your switch:

- [Generated Auto-QoS Configuration, page 27-18](#)
- [Effects of Auto-QoS on the Configuration, page 27-22](#)
- [Auto-QoS Configuration Guidelines, page 27-22](#)
- [Enabling Auto-QoS for VoIP, page 27-23](#)
- [Auto-QoS Configuration Example, page 27-24](#)

Generated Auto-QoS Configuration

By default, auto-QoS is disabled on all interfaces.

When auto-QoS is enabled, it uses the ingress packet label to categorize traffic, to assign packet labels, and to configure the ingress and egress queues as shown in [Table 27-2](#).

Table 27-2 Traffic Types, Ingress Packet Labels, Assigned Packet Labels, and Queues

	VoIP Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	All Other Traffic
Ingress DSCP	46	26	—	—	—
Ingress CoS	5	3	6	7	—
DiffServ	EF	AF31	—	—	—

Table 27-2 Traffic Types, Ingress Packet Labels, Assigned Packet Labels, and Queues (continued)

	VoIP Data Traffic	VoIP Control Traffic	Routing Protocol Traffic	STP BPDU Traffic	All Other Traffic
Assigned DSCP	46	26	48	56	0
Assigned CoS	5	3	6	7	0
CoS-to-Ingress Queue Map	2, 3, 4, 5, 6, 7 (queue 2)				0, 1 (queue 1)
CoS-to-Egress Queue Map	5 (queue 1)	3, 6, 7 (queue 2)			2, 4 (queue 3) 0, 1 (queue 4)

Table 27-3 shows the generated auto-QoS configuration for the ingress queues.

Table 27-3 Auto-QoS Configuration for the Ingress Queues

Ingress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size
SRR shared	1	0, 1	90 percent	90 percent
Priority	2	2, 3, 4, 5, 6, 7	10 percent	10 percent

Table 27-4 shows the generated auto-QoS configuration for the egress queues.

Table 27-4 Auto-QoS Configuration for the Egress Queues

Egress Queue	Queue Number	CoS-to-Queue Map	Queue Weight (Bandwidth)	Queue (Buffer) Size
Priority (shaped)	1	5	10 percent	20 percent
SRR shared	2	3, 6, 7	10 percent	20 percent
SRR shared	3	2, 4	60 percent	20 percent
SRR shared	4	0, 1	20 percent	40 percent

When you enable the auto-QoS feature on the first interface, these automatic actions occur:

- QoS is globally enabled (**mls qos** global configuration command), and other global configuration commands are added.
- When you enter the **auto qos voip cisco-phone** interface configuration command on a port at the edge of the network that is connected to a Cisco IP Phone, the switch enables the trusted boundary feature. The switch uses the Cisco Discovery Protocol (CDP) to detect the presence or absence of a Cisco IP Phone. When a Cisco IP Phone is detected, the ingress classification on the interface is set to trust the QoS label received in the packet. When a Cisco IP Phone is absent, the ingress classification is set to not trust the QoS label in the packet. The switch configures ingress and egress queues on the interface according to the settings in Table 27-3 and Table 27-4.

- When you enter the **auto qos voip trust** interface configuration command on a port connected to the interior of the network, the switch trusts the CoS value for nonrouted interfaces in ingress packets (the assumption is that traffic has already been classified by other edge devices). The switch configures the ingress and egress queues on the interface according to the settings in [Table 27-3](#) and [Table 27-4](#).

For information about the trusted boundary feature, see the “[Configuring a Trusted Boundary to Ensure Port Security](#)” section on page 27-34.

When you enable auto-QoS by using the **auto qos voip cisco-phone** or the **auto qos voip trust** interface configuration command, the switch automatically generates a QoS configuration based on the traffic type and ingress packet label and applies the commands listed in [Table 27-5](#) to the interface.

Table 27-5 Generated Auto-QoS Configuration

Description	Automatically Generated Command
The switch automatically enables standard QoS and configures the CoS-to-DSCP map (maps CoS values in incoming packets to a DSCP value).	<pre>Switch(config)# mls qos Switch(config)# mls qos map cos-dscp 0 8 16 26 32 46 48 56</pre>
The switch automatically maps CoS values to an ingress queue and to a threshold ID.	<pre>Switch(config)# no mls qos srr-queue input cos-map Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 3 0 Switch(config)# mls qos srr-queue input cos-map queue 1 threshold 2 1 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 1 2 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 2 4 6 7 Switch(config)# mls qos srr-queue input cos-map queue 2 threshold 3 3 5</pre>
The switch automatically maps CoS values to an egress queue and to a threshold ID.	<pre>Switch(config)# no mls qos srr-queue output cos-map Switch(config)# mls qos srr-queue output cos-map queue 1 threshold 3 5 Switch(config)# mls qos srr-queue output cos-map queue 2 threshold 3 3 6 7 Switch(config)# mls qos srr-queue output cos-map queue 3 threshold 3 2 4 Switch(config)# mls qos srr-queue output cos-map queue 4 threshold 2 1 Switch(config)# mls qos srr-queue output cos-map queue 4 threshold 3 0</pre>

Table 27-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command
The switch automatically maps DSCP values to an ingress queue and to a threshold ID.	<pre> Switch(config)# no mls qos srr-queue input dscp-map Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 9 10 11 12 13 14 15 Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 0 1 2 3 4 5 6 7 Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 3 32 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 1 16 17 18 19 20 21 22 23 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 26 33 34 35 36 37 38 39 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 48 49 50 51 52 53 54 55 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 2 56 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 24 25 27 28 29 30 31 40 Switch(config)# mls qos srr-queue input dscp-map queue 2 threshold 3 41 42 43 44 45 46 47 </pre>
The switch automatically maps DSCP values to an egress queue and to a threshold ID.	<pre> Switch(config)# no mls qos srr-queue output dscp-map Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 3 40 41 42 43 44 45 46 47 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 24 25 26 27 28 29 30 31 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 48 49 50 51 52 53 54 55 Switch(config)# mls qos srr-queue output dscp-map queue 2 threshold 3 56 57 58 59 60 61 62 63 Switch(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 16 17 18 19 20 21 22 23 Switch(config)# mls qos srr-queue output dscp-map queue 3 threshold 3 32 33 34 35 36 37 38 39 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 1 8 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 2 9 10 11 12 13 14 15 Switch(config)# mls qos srr-queue output dscp-map queue 4 threshold 3 0 1 2 3 4 5 6 7 </pre>
The switch automatically sets up the ingress queues, with queue 2 as the priority queue and queue 1 in shared mode. The switch also configures the bandwidth and buffer size for the ingress queues.	<pre> Switch(config)# no mls qos srr-queue input priority-queue 1 Switch(config)# no mls qos srr-queue input priority-queue 2 Switch(config)# mls qos srr-queue input bandwidth 90 10 Switch(config)# no mls qos srr-queue input buffers </pre>
The switch automatically configures the egress queue buffer sizes. It configures the bandwidth and the SRR mode (shaped or shared) on the egress queues mapped to the port.	<pre> Switch(config)# mls qos queue-set output 1 buffers 20 20 20 40 Switch(config-if)# srr-queue bandwidth shape 10 0 0 0 Switch(config-if)# srr-queue bandwidth share 10 10 60 20 </pre>

Table 27-5 Generated Auto-QoS Configuration (continued)

Description	Automatically Generated Command
The switch automatically sets the ingress classification to trust the CoS value received in the packet on a nonrouted interface.	Switch(config-if)# mls qos trust cos Switch(config-if)# mls qos trust dscp
If you entered the auto qos voip cisco-phone command, the switch automatically enables the trusted boundary feature, which uses the CDP to detect the presence or absence of a Cisco IP Phone.	Switch(config-if)# mls qos trust device cisco-phone

Effects of Auto-QoS on the Configuration

When auto-QoS is enabled, the **auto qos voip** interface configuration command and the generated configuration are added to the running configuration.

The switch applies the auto-QoS-generated commands as if the commands were entered from the CLI. An existing user configuration can cause the application of the generated commands to fail or to be overridden by the generated commands. These actions occur without warning. If all the generated commands are successfully applied, any user-entered configuration that was not overridden remains in the running configuration. Any user-entered configuration that was overridden can be retrieved by reloading the switch without saving the current configuration to memory. If the generated commands fail to be applied, the previous running configuration is restored.

Auto-QoS Configuration Guidelines

Before configuring auto-QoS, you should be aware of this information:

- In this release, auto-QoS configures the switch only for VoIP with Cisco IP Phones.
- To take advantage of the auto-QoS defaults, you should enable auto-QoS before you configure other QoS commands. If necessary, you can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed. For more information, see the [“Effects of Auto-QoS on the Configuration” section on page 27-22](#).
- You can enable auto-QoS on static, dynamic-access, voice VLAN access, and trunk ports.
- By default, the CDP is enabled on all interfaces. For auto-QoS to function properly, do not disable the CDP.
- Policing is not enabled with auto-QoS. You can manually enable policing, as described in the [“Configuring a QoS Policy” section on page 27-36](#).

Enabling Auto-QoS for VoIP

Beginning in privileged EXEC mode, follow these steps to enable auto-QoS for VoIP within a QoS domain:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface that is connected to a Cisco IP Phone or the uplink interface that is connected to another switch or router in the interior of the network.
Step 3	auto qos voip { cisco-phone trust }	<p>Enable auto-QoS.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • cisco-phone—If the interface is connected to a Cisco IP Phone, the QoS labels of incoming packets are trusted only when the telephone is detected. • trust—The uplink interface is connected to a trusted switch or router, and the VoIP traffic classification in the ingress packet is trusted.
Step 4	end	Return to privileged EXEC mode.
Step 5	show auto qos interface <i>interface-id</i>	<p>Verify your entries.</p> <p>This command displays the initial auto-QoS configuration that was applied; it does not display any user changes to the configuration that might be in effect. You can use the show running-config privileged EXEC command to display the auto-QoS configuration and the user modifications.</p>

To display the QoS commands that are automatically generated when auto-QoS is enabled or disabled, enter the **debug autoqos** privileged EXEC command *before* enabling auto-QoS. For more information, refer to the “debug autoqos” command in the command reference for this release.

To disable auto-QoS on an interface, use the **no auto qos voip** interface configuration command. Only the auto-QoS-generated interface configuration commands for this interface are removed. If this is the last interface on which auto-QoS is enabled and you enter the **no auto qos voip** command, auto-QoS is considered disabled even though the auto-QoS-generated global configuration commands remain (to avoid disrupting traffic on other interfaces affected by the global configuration). You can use the **no mls qos** global configuration command to disable the auto-QoS-generated global configuration commands. With QoS disabled, there is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

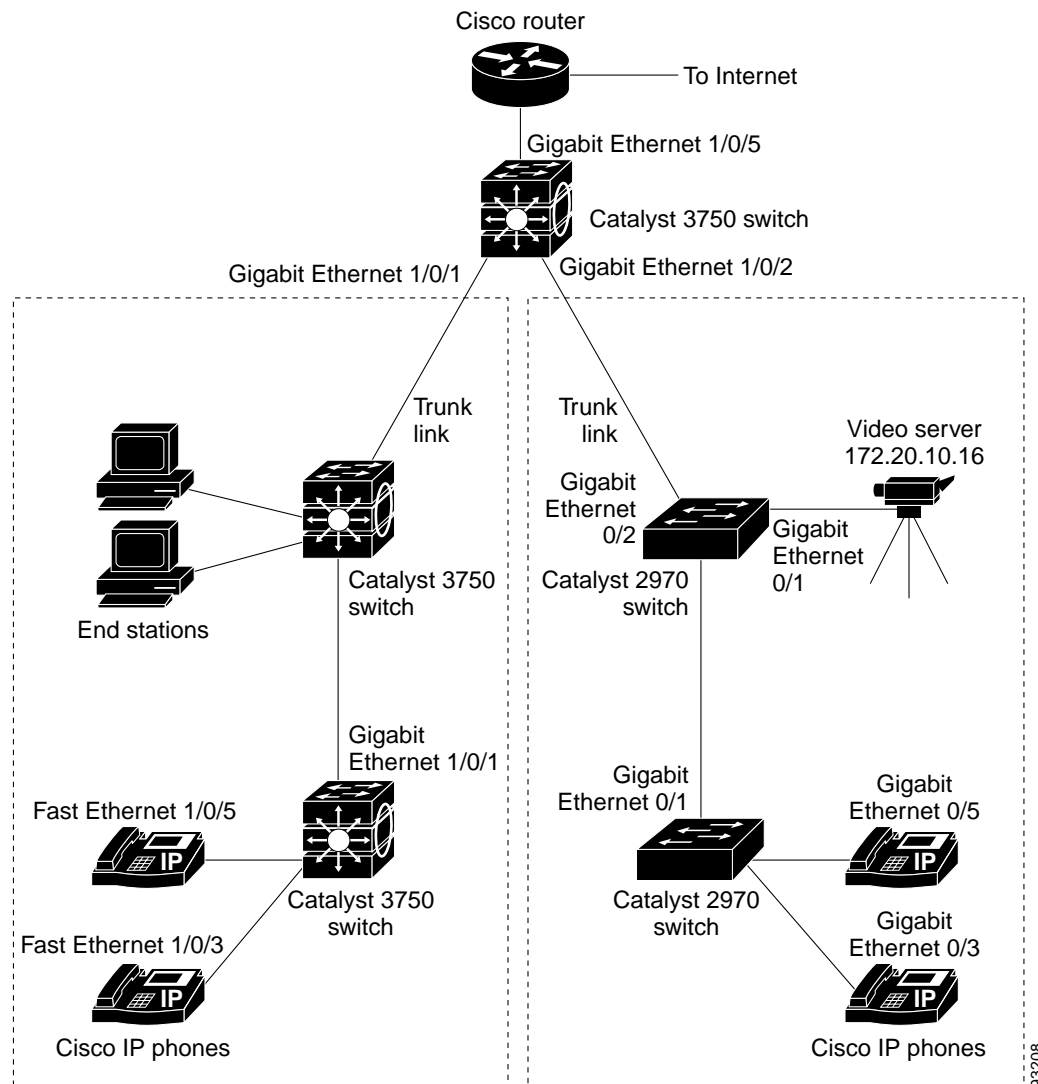
This example shows how to enable auto-QoS and to trust the QoS labels received in incoming packets when the switch or router connected to Gigabit Ethernet interface 0/1 is a trusted device:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# auto qos voip trust
```

Auto-QoS Configuration Example

This section describes how you could implement auto-QoS in a network, as shown in [Figure 27-10](#).

Figure 27-10 Auto-QoS Configuration Example Network



The intelligent wiring closets in [Figure 27-10](#) contain Catalyst 2970 and Catalyst 3750 switches. The object of this example is to prioritize the VoIP traffic over all other traffic. To do so, enable auto-QoS on the switches at the edge of the QoS domains in the wiring closets.

**Note**

You should not configure any standard QoS commands before entering the auto-QoS commands. You can fine-tune the QoS configuration, but we recommend that you do so only after the auto-QoS configuration is completed.

Beginning in privileged EXEC mode, follow these steps to configure the switch at the edge of the QoS domain to prioritize the VoIP traffic over all other traffic:

	Command	Purpose
Step 1	debug autoqos	Enable debugging for auto-QoS. When debugging is enabled, the switch displays the QoS configuration that is automatically generated when auto-QoS is enabled.
Step 2	configure terminal	Enter global configuration mode.
Step 3	cdp enable	Enable CDP globally. By default, it is enabled.
Step 4	interface gigabitethernet0/3	Enter interface configuration mode.
Step 5	auto qos voip cisco-phone	Enable auto-QoS on the interface, and specify that the interface is connected to a Cisco IP Phone. The QoS labels of incoming packets are trusted only when the Cisco IP Phone is detected.
Step 6	exit	Return to global configuration mode.
Step 7	interface gigabitethernet0/4	Enter interface configuration mode, and specify that the interface is connected to a Cisco IP Phone. The QoS labels of incoming packets are trusted only when the Cisco IP Phone is detected.
Step 8	auto qos voip cisco-phone	Enable auto-QoS on the interface, and specify that the interface is connected to a Cisco IP Phone.
Step 9	exit	Return to global configuration mode.
Step 10	interface gigabitethernet0/1	Enter interface configuration mode.
Step 11	auto qos voip trust	Enable auto-QoS on the interface, and specify that the interface is connected to a trusted router or switch.
Step 12	end	Return to privileged EXEC mode.
Step 13	show auto qos	Verify your entries. This command displays the auto-QoS configuration that is initially applied; it does not display any user changes to the configuration that might be in effect. For information about the QoS configuration that might be affected by auto-QoS, see the “Displaying Auto-QoS Information” section on page 26-12.
Step 14	copy running-config startup-config	Save the auto qos voip interface configuration commands and the generated auto-QoS configuration in the configuration file.

Displaying Auto-QoS Information

To display the initial auto-QoS configuration, use the **show auto qos [interface *interface-id*]** privileged EXEC command. To display any user changes to that configuration, use the **show running-config** privileged EXEC command. You can compare the **show auto qos** and the **show running-config** command output to identify the user-defined QoS settings.

To display information about the QoS configuration that might be affected by auto-QoS, use one of these commands:

- **show mls qos**
- **show mls qos maps cos-dscp**
- **show mls qos interface [interface-id] [buffers | queueing]**
- **show mls qos maps [cos-dscp | cos-input-q | cos-output-q | dscp-cos | dscp-input-q | dscp-output-q]**
- **show mls qos input-queue**
- **show running-config**

For more information about these commands, refer to the command reference for this release.

Configuring Standard QoS

Before configuring standard QoS, you must have a thorough understanding of these items:

- The types of applications used and the traffic patterns on your network.
- Traffic characteristics and needs of your network. Is the traffic bursty? Do you need to reserve bandwidth for voice and video streams?
- Bandwidth requirements and speed of the network.
- Location of congestion points in the network.

These sections describe how to configure QoS on your switch:

- [Default Standard QoS Configuration, page 27-27](#)
- [Standard QoS Configuration Guidelines, page 27-29](#)
- [Enabling QoS Globally, page 27-30](#) (required)
- [Configuring Classification Using Port Trust States, page 27-30](#) (required)
- [Configuring a QoS Policy, page 27-36](#) (required)
- [Configuring DSCP Maps, page 27-47](#) (optional, unless you need to use the DSCP-to-DSCP-mutation map or the policed-DSCP map)
- [Configuring Ingress Queue Characteristics, page 27-52](#) (optional)
- [Configuring Egress Queue Characteristics, page 27-57](#) (optional)

Default Standard QoS Configuration

QoS is disabled. There is no concept of trusted or untrusted ports because the packets are not modified (the CoS, DSCP, and IP precedence values in the packet are not changed). Traffic is switched in pass-through mode (packets are switched without any rewrites and classified as best effort without any policing).

When QoS is enabled with the **mls qos** global configuration command and all other QoS settings are at their defaults, traffic is classified as best effort (the DSCP and CoS value is set to 0) without any policing. No policy maps are configured. The default port trust state on all ports is untrusted. The default ingress and egress queue settings are described in the [“Default Ingress Queue Configuration” section on page 27-27](#) and the [“Default Egress Queue Configuration” section on page 27-28](#).

Default Ingress Queue Configuration

[Table 27-6](#) shows the default ingress queue configuration when QoS is enabled.

Table 27-6 Default Ingress Queue Configuration

Feature	Queue 1	Queue 2
Buffer Allocation	90 percent	10 percent
Bandwidth Allocation ¹	4	4
Priority Queue Bandwidth ²	0	10
WTD Drop Threshold 1	100 percent	100 percent
WTD Drop Threshold 2	100 percent	100 percent

1. The bandwidth is equally shared between the queues. SRR sends packets in shared mode only.
2. Queue 2 is the priority queue. SRR services the priority queue for its configured share before servicing the other queue.

[Table 27-7](#) shows the default CoS input queue threshold map when QoS is enabled.

Table 27-7 Default CoS Input Queue Threshold Map

CoS Value	0–4	5	6, 7
Queue ID - Threshold ID	1 - 1	2 - 1	1 - 1

[Table 27-8](#) shows the default DSCP input queue threshold map when QoS is enabled.

Table 27-8 Default DSCP Input Queue Threshold Map

DSCP Value	0–39	40–47	48–63
Queue ID - Threshold ID	1 - 1	2 - 1	1 - 1

Default Egress Queue Configuration

[Table 27-9](#) shows the default egress queue configuration for each queue-set when QoS is enabled. All ports are mapped to queue-set 1. The port bandwidth limit is set to 100 percent and rate unlimited.

Table 27-9 Default Egress Queue Configuration

Feature	Queue 1	Queue 2	Queue 3	Queue 4
Buffer Allocation	25 percent	25 percent	25 percent	25 percent
WTD Drop Threshold 1	100 percent	50 percent	100 percent	100 percent
WTD Drop Threshold 2	100 percent	50 percent	100 percent	100 percent
Reserved Threshold	50 percent	100 percent	50 percent	50 percent
Maximum Threshold	400 percent	400 percent	400 percent	400 percent
SRR Shaped Weights (absolute) ¹	25	0	0	0
SRR Shared Weights ²	25	25	25	25

1. A shaped weight of zero means that this queue is operating in shared mode.
2. One quarter of the bandwidth is allocated to each queue.

[Table 27-10](#) shows the default CoS output queue threshold map when QoS is enabled.

Table 27-10 Default CoS Output Queue Threshold Map

CoS Value	0, 1	2, 3	4	5	6, 7
Queue ID - Threshold ID	2 - 1	3 - 1	4 - 1	1 - 1	4 - 1

[Table 27-11](#) shows the default DSCP output queue threshold map when QoS is enabled.

Table 27-11 Default DSCP Output Queue Threshold Map

DSCP Value	0–15	16–31	32–39	40–47	48–63
Queue ID - Threshold ID	2 - 1	3 - 1	4 - 1	1 - 1	4 - 1

Default Mapping Table Configuration

The default CoS-to-DSCP map is shown in [Table 27-12 on page 27-47](#).

The default IP-precedence-to-DSCP map is shown in [Table 27-13 on page 27-48](#).

The default DSCP-to-CoS map is shown in [Table 27-14 on page 27-50](#).

The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

The default policed-DSCP map is a null map, which maps an incoming DSCP value to the same DSCP value (no markdown).

Standard QoS Configuration Guidelines

Before beginning the QoS configuration, you should be aware of this information:

- You configure QoS only on physical ports; there is no support for it on the VLAN or switch virtual interface level.
- It is not possible to match IP fragments against configured IP extended ACLs to enforce QoS. IP fragments are sent as best-effort. IP fragments are denoted by fields in the IP header.
- Only one ACL per class map and only one **match** class-map configuration command per class map are supported. The ACL can have multiple ACEs, which match fields against the contents of the packet.
- Incoming traffic is classified, policed, and marked down (if configured) regardless of whether the traffic is bridged, routed, or sent to the CPU. It is possible for bridged frames to be dropped or to have their DSCP and CoS values modified.
- Only one policer is applied to a packet on an ingress interface. Only the average rate and committed burst parameters are configurable.
- The port ASIC device, which controls more than one physical port, supports 256 policers (255 policers plus 1 **no** policer). The maximum number of policers supported per port is 64. Policers are allocated on demand by the software and are constrained by the hardware and ASIC boundaries. You cannot reserve policers per port; there is no guarantee that a port will be assigned to any policer.
- On an interface configured for QoS, all traffic received through the interface is classified, policed, and marked according to the policy map attached to the interface. On a trunk interface configured for QoS, traffic in *all* VLANs received through the interface is classified, policed, and marked according to the policy map attached to the interface.
- You can create an aggregate policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps.
- If you have EtherChannel ports configured on your switch, you must configure QoS classification, policing, mapping, and queueing on the individual physical ports that comprise the EtherChannel. You must decide whether the QoS configuration should match on all ports in the EtherChannel.
- Control traffic (such as spanning-tree bridge protocol data units [BPDUs] and routing update packets) received by the switch are subject to all ingress QoS processing.
- You are likely to lose data when you change queue settings; therefore, try to make changes when traffic is at a minimum.

Enabling QoS Globally

By default, QoS is disabled on the switch.

Beginning in privileged EXEC mode, follow these steps to enable QoS. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS globally. QoS runs from the default settings described in the “Default Standard QoS Configuration” section on page 27-27, the “Queueing and Scheduling on Ingress Queues” section on page 27-13, and the “Queueing and Scheduling on Egress Queues” section on page 27-15.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable QoS, use the **no mls qos** global configuration command.

Configuring Classification Using Port Trust States

These sections describe how to classify incoming traffic by using port trust states. Depending on your network configuration, you must perform one or more of these tasks or one or more of the tasks in the [“Configuring a QoS Policy”](#) section on page 27-36:

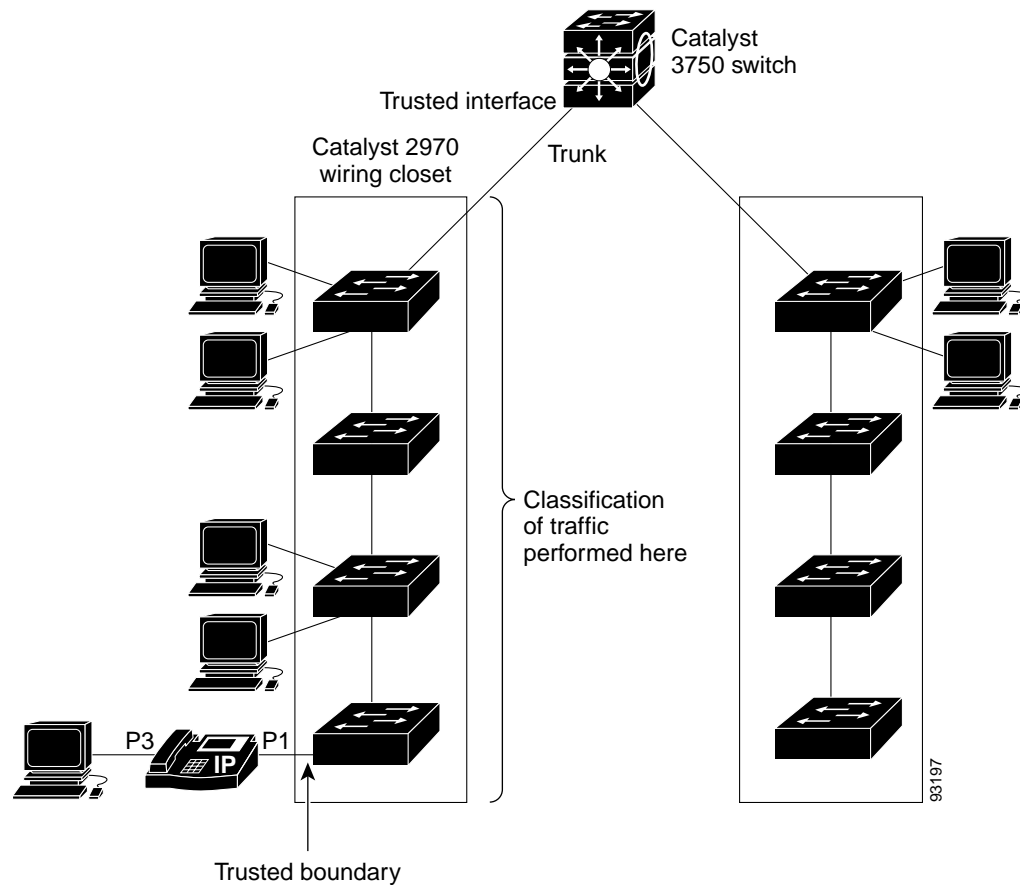
- [Configuring the Trust State on Ports within the QoS Domain](#), page 27-31
- [Configuring the CoS Value for an Interface](#), page 27-33
- [Configuring a Trusted Boundary to Ensure Port Security](#), page 27-34
- [Configuring the DSCP Trust State on a Port Bordering Another QoS Domain](#), page 27-35

Configuring the Trust State on Ports within the QoS Domain

Packets entering a QoS domain are classified at the edge of the QoS domain. When the packets are classified at the edge, the switch port within the QoS domain can be configured to one of the trusted states because there is no need to classify the packets at every switch within the QoS domain.

Figure 27-11 shows a sample network topology.

Figure 27-11 Port Trusted States within the QoS Domain



Beginning in privileged EXEC mode, follow these steps to configure the port to trust the classification of the traffic that it receives:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 3	mls qos trust [cos dscp ip-precedence]	Configure the port trust state. By default, the port is not trusted. If no keyword is specified, the default is dscp . The keywords have these meanings: <ul style="list-style-type: none"> • cos—Classifies an ingress packet by using the packet CoS value. For an untagged packet, the port default CoS value is used. The default port CoS value is 0. • dscp—Classifies an ingress packet by using the packet DSCP value. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map. • ip-precedence—Classifies an ingress packet by using the packet IP-precedence value. For a non-IP packet, the packet CoS value is used if the packet is tagged; for an untagged packet, the default port CoS is used. Internally, the switch maps the CoS value to a DSCP value by using the CoS-to-DSCP map.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return a port to its untrusted state, use the **no mls qos trust** interface configuration command.

For information on how to change the default CoS value, see the [“Configuring the CoS Value for an Interface” section on page 27-33](#). For information on how to configure the CoS-to-DSCP map, see the [“Configuring the CoS-to-DSCP Map” section on page 27-47](#).

Configuring the CoS Value for an Interface

QoS assigns the CoS value specified with the **mls qos cos** interface configuration command to untagged frames received on trusted and untrusted ports.

Beginning in privileged EXEC mode, follow these steps to define the default CoS value of a port or to assign the default CoS to all incoming packets on the port:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured. Valid interfaces include physical interfaces.
Step 3	mls qos cos { <i>default-cos</i> override }	Configure the default CoS value for the port. <ul style="list-style-type: none"> For <i>default-cos</i>, specify a default CoS value to be assigned to a port. If the packet is untagged, the default CoS value becomes the packet CoS value. The CoS range is 0 to 7. The default is 0. Use the override keyword to override the previously configured trust state of the incoming packet and to apply the default port CoS value to the port on all incoming packets. By default, CoS override is disabled. <p>Use the override keyword when all incoming packets on specified ports deserve higher or lower priority than packets entering from other ports. Even if a port was previously set to trust DSCP, CoS, or IP precedence, this command overrides the previously configured trust state, and all the incoming CoS values are assigned the default CoS value configured with this command. If an incoming packet is tagged, the CoS value of the packet is modified with the default CoS of the port at the ingress port.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos cos** {*default-cos* | **override**} interface configuration command.

Configuring a Trusted Boundary to Ensure Port Security

In a typical network, you connect a Cisco IP Phone to a switch port, as shown in [Figure 27-11 on page 27-31](#), and cascade devices that generate data packets from the back of the telephone. The Cisco IP Phone guarantees the voice quality through a shared data link by marking the CoS level of the voice packets as high priority (CoS = 5) and by marking the data packets as low priority (CoS = 0). Traffic sent from the telephone to the switch is typically marked with a tag that uses the 802.1Q header. The header contains the VLAN information and the class of service (CoS) 3-bit field, which determines the priority of the packet.

For most Cisco IP Phone configurations, the traffic sent from the telephone to the switch should be trusted to ensure that voice traffic is properly prioritized over other types of traffic in the network. By using the **mls qos trust cos** interface configuration command, you configure the switch port to which the telephone is connected to trust the CoS labels of all traffic received on that port.

With the trusted setting, you also can use the trusted boundary feature to prevent misuse of a high-priority queue if a user bypasses the telephone and connects the PC directly to the switch. Without trusted boundary, the CoS labels generated by the PC are trusted by the switch (because of the trusted CoS setting). By contrast, trusted boundary uses CDP to detect the presence of a Cisco IP Phone (such as the Cisco IP Phone 7910, 7935, 7940, and 7960) on a switch port. If the telephone is not detected, the trusted boundary feature disables the trusted setting on the switch port and prevents misuse of a high-priority queue. Note that the trusted boundary feature is not effective if the PC and Cisco IP Phone are connected to a hub that is connected to the switch.

In some situations, you can prevent a PC connected to the Cisco IP Phone from taking advantage of a high-priority data queue. You can use the **switchport priority extend cos** interface configuration command to configure the telephone through the switch CLI to override the priority of the traffic received from the PC.

Beginning in privileged EXEC mode, follow these steps to enable trusted boundary on a port:

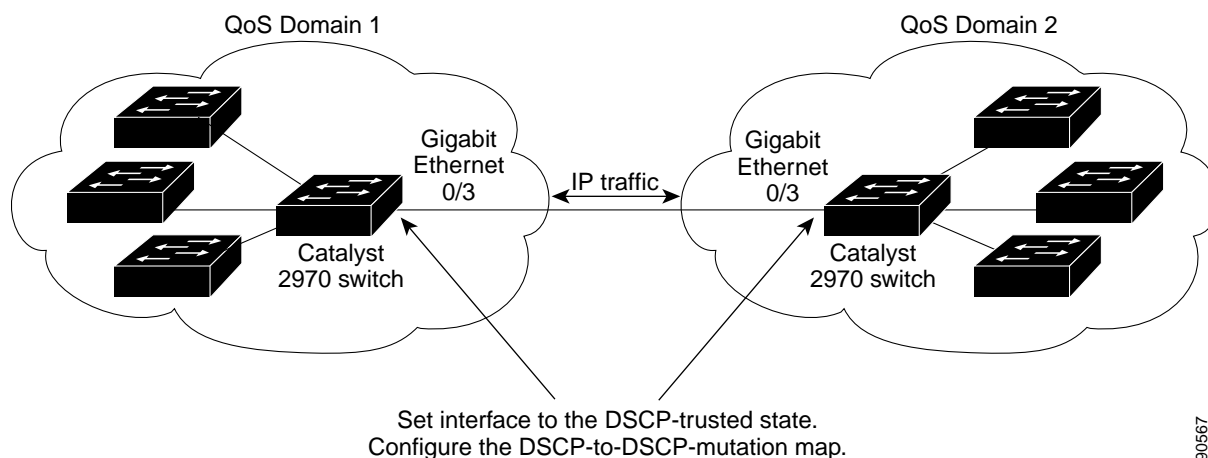
	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cdp run	Enable CDP globally. By default, CDP is enabled.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface connected to the Cisco IP Phone. Valid interfaces include physical interfaces.
Step 4	cdp enable	Enable CDP on the interface. By default, CDP is enabled.
Step 5	mls qos trust cos	Configure the interface to trust the CoS value in traffic received from the Cisco IP Phone. By default, the port is not trusted.
Step 6	mls qos trust device cisco-phone	Specify that the Cisco IP Phone is a trusted device. You cannot enable both trusted boundary and auto-QoS (auto qos voip interface configuration command) at the same time; they are mutually exclusive.
Step 7	end	Return to privileged EXEC mode.
Step 8	show mls qos interface	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the trusted boundary feature, use the **no mls qos trust device** interface configuration command.

Configuring the DSCP Trust State on a Port Bordering Another QoS Domain

If you are administering two separate QoS domains between which you want to implement QoS features for IP traffic, you can configure the switch ports bordering the domains to a DSCP-trusted state as shown in Figure 27-12. Then the receiving port accepts the DSCP-trusted value and avoids the classification stage of QoS. If the two domains use different DSCP values, you can configure the DSCP-to-DSCP-mutation map to translate a set of DSCP values to match the definition in the other domain.

Figure 27-12 DSCP-Trusted State on a Port Bordering Another QoS Domain



Beginning in privileged EXEC mode, follow these steps to configure the DSCP-trusted state on a port and modify the DSCP-to-DSCP-mutation map. To ensure a consistent mapping strategy across both QoS domains, you must perform this procedure on the ports in both domains:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map dscp-mutation <i>dscp-mutation-name in-dscp to out-dscp</i>	Modify the DSCP-to-DSCP-mutation map. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value. <ul style="list-style-type: none"> For <i>dscp-mutation-name</i>, enter the mutation map name. You can create more than one map by specifying a new name. For <i>in-dscp</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. For <i>out-dscp</i>, enter a single DSCP value. The DSCP range is 0 to 63.
Step 3	interface interface-id	Enter interface configuration mode, and specify the interface to be trusted. Valid interfaces include physical interfaces.
Step 4	mls qos trust dscp	Configure the ingress port as a DSCP-trusted port. By default, the port is not trusted.

	Command	Purpose
Step 5	mls qos dscp-mutation <i>dscp-mutation-name</i>	Apply the map to the specified ingress DSCP-trusted port. For <i>dscp-mutation-name</i> , specify the mutation map name created in Step 2. You can configure multiple DSCP-to-DSCP-mutation maps on an ingress port.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos maps dscp-mutation	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return a port to its non-trusted state, use the **no mls qos trust** interface configuration command. To return to the default DSCP-to-DSCP-mutation map values, use the **no mls qos map dscp-mutation** *dscp-mutation-name* global configuration command.

This example shows how to configure Gigabit Ethernet port 0/3 to the DSCP-trusted state and to modify the DSCP-to-DSCP-mutation map (named *gi0/3-mutation*) so that incoming DSCP values 10 to 13 are mapped to DSCP 30:

```
Switch(config)# mls qos map dscp-mutation gi0/3-mutation 10 11 12 13 to 30
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation gi0/3-mutation
Switch(config-if)# end
```

Configuring a QoS Policy

Configuring a QoS policy typically requires classifying traffic into classes, configuring policies applied to those traffic classes, and attaching policies to interfaces.

For background information, see the [“Classification” section on page 27-4](#) and the [“Policing and Marking” section on page 27-8](#). For configuration guidelines, see the [“Standard QoS Configuration Guidelines” section on page 27-29](#).

These sections describe how to classify, police, and mark traffic. Depending on your network configuration, you must perform one or more of these tasks:

- [Classifying Traffic by Using ACLs, page 27-37](#)
- [Classifying Traffic by Using Class Maps, page 27-40](#)
- [Classifying, Policing, and Marking Traffic by Using Policy Maps, page 27-42](#)
- [Classifying, Policing, and Marking Traffic by Using Aggregate Policers, page 27-45](#)

Classifying Traffic by Using ACLs

You can classify IP traffic by using IP standard or IP extended ACLs; you can classify non-IP traffic by using Layer 2 MAC ACLs.

Beginning in privileged EXEC mode, follow these steps to create an IP standard ACL for IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>]	<p>Create an IP standard ACL, repeating the command as many times as necessary.</p> <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number. The range is 1 to 99 and 1300 to 1999. Use the permit keyword to permit a certain type of traffic if the conditions are matched. Use the deny keyword to deny a certain type of traffic if conditions are matched. For <i>source</i>, enter the network or host from which the packet is being sent. You can use the any keyword as an abbreviation for 0.0.0.0 255.255.255.255. (Optional) For <i>source-wildcard</i>, enter the wildcard bits in dotted decimal notation to be applied to the source. Place ones in the bit positions that you want to ignore. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to allow access for only those hosts on the three specified networks. The wildcard bits apply to the host portions of the network addresses. Any host with a source address that does not match the access list statements is rejected.

```
Switch(config)# access-list 1 permit 192.5.255.0 0.0.0.255
Switch(config)# access-list 1 permit 128.88.0.0 0.0.255.255
Switch(config)# access-list 1 permit 36.0.0.0 0.0.0.255
! (Note: all other access implicitly denied)
```

Beginning in privileged EXEC mode, follow these steps to create an IP extended ACL for IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i>	<p>Create an IP extended ACL, repeating the command as many times as necessary.</p> <ul style="list-style-type: none"> For <i>access-list-number</i>, enter the access list number. The range is 100 to 199 and 2000 to 2699. Use the permit keyword to permit a certain type of traffic if the conditions are matched. Use the deny keyword to deny a certain type of traffic if conditions are matched. For <i>protocol</i>, enter the name or number of an IP protocol. Use the question mark (?) to see a list of available protocol keywords. For <i>source</i>, enter the network or host from which the packet is being sent. You specify this by using dotted decimal notation, by using the any keyword as an abbreviation for <i>source</i> 0.0.0.0 <i>source-wildcard</i> 255.255.255.255, or by using the host keyword for <i>source</i> 0.0.0.0. For <i>source-wildcard</i>, enter the wildcard bits by placing ones in the bit positions that you want to ignore. You specify the wildcard by using dotted decimal notation, by using the any keyword as an abbreviation for <i>source</i> 0.0.0.0 <i>source-wildcard</i> 255.255.255.255, or by using the host keyword for <i>source</i> 0.0.0.0. For <i>destination</i>, enter the network or host to which the packet is being sent. You have the same options for specifying the <i>destination</i> and <i>destination-wildcard</i> as those described by <i>source</i> and <i>source-wildcard</i>. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
Step 3	end	Return to privileged EXEC mode.
Step 4	show access-lists	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no access-list** *access-list-number* global configuration command.

This example shows how to create an ACL that permits IP traffic from any source to any destination that has the DSCP value set to 32:

```
Switch(config)# access-list 100 permit ip any any dscp 32
```

This example shows how to create an ACL that permits IP traffic from a source host at 10.1.1.1 to a destination host at 10.1.1.2 with a precedence value of 5:

```
Switch(config)# access-list 100 permit ip host 10.1.1.1 host 10.1.1.2 precedence 5
```

This example shows how to create an ACL that permits PIM traffic from any source to a destination group address of 224.0.0.2 with a DSCP set to 32:

```
Switch(config)# access-list 102 permit pim any 224.0.0.2 dscp 32
```

Beginning in privileged EXEC mode, follow these steps to create a Layer 2 MAC ACL for non-IP traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mac access-list extended <i>name</i>	Create a Layer 2 MAC ACL by specifying the name of the list. After entering this command, the mode changes to extended MAC ACL configuration.
Step 3	{permit deny} {host <i>src-MAC-addr mask</i> any host <i>dst-MAC-addr</i> <i>dst-MAC-addr mask</i> } [<i>type mask</i>]	Specify the type of traffic to permit or deny if the conditions are matched, entering the command as many times as necessary. <ul style="list-style-type: none"> For <i>src-MAC-addr</i>, enter the MAC address of the host from which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the any keyword as an abbreviation for <i>source</i> 0.0.0, <i>source-wildcard</i> 255.255.255, or by using the host keyword for <i>source</i> 0.0.0. For <i>mask</i>, enter the wildcard bits by placing ones in the bit positions that you want to ignore. For <i>dst-MAC-addr</i>, enter the MAC address of the host to which the packet is being sent. You specify this by using the hexadecimal format (H.H.H), by using the any keyword as an abbreviation for <i>source</i> 0.0.0, <i>source-wildcard</i> 255.255.255, or by using the host keyword for <i>source</i> 0.0.0. (Optional) For <i>type mask</i>, specify the Ethertype number of a packet with Ethernet II or SNAP encapsulation to identify the protocol of the packet. For <i>type</i>, the range is from 0 to 65535, typically specified in hexadecimal. For <i>mask</i>, enter the <i>don't care</i> bits applied to the Ethertype before testing for a match. <p>Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show access-lists [<i>access-list-number</i> <i>access-list-name</i>]	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an access list, use the **no mac access-list extended** *access-list-name* global configuration command.

This example shows how to create a Layer 2 MAC ACL with two permit statements. The first statement allows traffic from the host with MAC address 0001.0000.0001 to the host with MAC address 0002.0000.0001. The second statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 to the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-macl)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-macl)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
! (Note: all other access implicitly denied)
```

Classifying Traffic by Using Class Maps

You use the **class-map** global configuration command to name and to isolate a specific traffic flow (or class) from all other traffic. The class map defines the criteria to use to match against a specific traffic flow to further classify it. Match statements can include criteria such as an ACL, IP precedence values, or DSCP values. The match criterion is defined with one match statement entered within the class-map configuration mode.



Note

You can also create class-maps during policy map creation by using the **class** policy-map configuration command. For more information, see the [“Classifying, Policing, and Marking Traffic by Using Policy Maps” section on page 27-42](#).

Beginning in privileged EXEC mode, follow these steps to create a class map and to define the match criterion to classify traffic:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	access-list <i>access-list-number</i> { deny permit } <i>source</i> [<i>source-wildcard</i>] or access-list <i>access-list-number</i> { deny permit } <i>protocol</i> <i>source</i> [<i>source-wildcard</i>] <i>destination</i> [<i>destination-wildcard</i>] or mac access-list extended <i>name</i> { permit deny } { host <i>src-MAC-addr</i> <i>mask</i> any host <i>dst-MAC-addr</i> <i>dst-MAC-addr</i> <i>mask</i> } [<i>type</i> <i>mask</i>]	Create an IP standard or extended ACL for IP traffic or a Layer 2 MAC ACL for non-IP traffic, repeating the command as many times as necessary. For more information, see the “Classifying Traffic by Using ACLs” section on page 27-37 . Note When creating an access list, remember that, by default, the end of the access list contains an implicit deny statement for everything if it did not find a match before reaching the end.
Step 3	class-map [match-all match-any] <i>class-map-name</i>	Create a class map, and enter class-map configuration mode. By default, no class maps are defined. <ul style="list-style-type: none"> (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. (Optional) Use the match-any keyword to perform a logical-OR of all matching statements under this class map. One or more match criteria must be matched. For <i>class-map-name</i>, specify the name of the class map. If neither the match-all or match-any keyword is specified, the default is match-all . Note Because only one match command per class map is supported, the match-all and match-any keywords function the same.

	Command	Purpose
Step 4	match { access-group <i>acl-index-or-name</i> ip dscp <i>dscp-list</i> ip precedence <i>ip-precedence-list</i> }	<p>Define the match criterion to classify traffic.</p> <p>By default, no match criterion is defined.</p> <p>Only one match criterion per class map is supported, and only one ACL per class map is supported.</p> <ul style="list-style-type: none"> For access-group <i>acl-index-or-name</i>, specify the number or name of the ACL created in Step 2. For ip dscp <i>dscp-list</i>, enter a list of up to eight IP DSCP values to match against incoming packets. Separate each value with a space. The range is 0 to 63. For ip precedence <i>ip-precedence-list</i>, enter a list of up to eight IP-precedence values to match against incoming packets. Separate each value with a space. The range is 0 to 7.
Step 5	end	Return to privileged EXEC mode.
Step 6	show class-map	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing class map, use the **no class-map** [**match-all** | **match-any**] *class-map-name* global configuration command. To remove a match criterion, use the **no match** {**access-group** *acl-index-or-name* | **ip dscp** | **ip precedence**} class-map configuration command.

This example shows how to configure the class map called *class1*. The *class1* has one match criterion, which is access list 103. It permits traffic from any host to any destination that matches a DSCP value of 10.

```
Switch(config)# access-list 103 permit any any dscp 10
Switch(config)# class-map class1
Switch(config-cmap)# match access-group 103
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class2*, which matches incoming traffic with DSCP values of 10, 11, and 12.

```
Switch(config)# class-map class2
Switch(config-cmap)# match ip dscp 10 11 12
Switch(config-cmap)# end
Switch#
```

This example shows how to create a class map called *class3*, which matches incoming traffic with IP-precedence values of 5, 6, and 7:

```
Switch(config)# class-map class3
Switch(config-cmap)# match ip precedence 5 6 7
Switch(config-cmap)# end
Switch#
```

Classifying, Policing, and Marking Traffic by Using Policy Maps

A policy map specifies which traffic class to act on. Actions can include trusting the CoS, DSCP, or IP precedence values in the traffic class; setting a specific DSCP or IP precedence value in the traffic class; and specifying the traffic bandwidth limitations for each matched traffic class (policer) and the action to take when the traffic is out of profile (marking).

A policy map also has these characteristics:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.
- A policy-map trust state and an interface trust state are mutually exclusive, and whichever is configured last takes affect.

You can attach only one policy map per ingress interface.

Beginning in privileged EXEC mode, follow these steps to create a policy map:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	class-map [match-all match-any] <i>class-map-name</i>	<p>Create a class map, and enter class-map configuration mode.</p> <p>By default, no class maps are defined.</p> <ul style="list-style-type: none"> • (Optional) Use the match-all keyword to perform a logical-AND of all matching statements under this class map. All match criteria in the class map must be matched. • (Optional) Use the match-any keyword to perform a logical-OR of all matching statements under this class map. One or more match criteria must be matched. • For <i>class-map-name</i>, specify the name of the class map. <p>If neither the match-all or match-any keyword is specified, the default is match-all.</p> <p>Note Because only one match command per class map is supported, the match-all and match-any keywords function the same.</p>
Step 3	policy-map <i>policy-map-name</i>	<p>Create a policy map by entering the policy map name, and enter policy-map configuration mode.</p> <p>By default, no policy maps are defined.</p> <p>The default behavior of a policy map is to set the DSCP to 0 if the packet is an IP packet and to set the CoS to 0 if the packet is tagged. No policing is performed.</p>
Step 4	class <i>class-map-name</i>	<p>Define a traffic classification, and enter policy-map class configuration mode.</p> <p>By default, no policy map class-maps are defined.</p> <p>If a traffic class has already been defined by using the class-map global configuration command, specify its name for <i>class-map-name</i> in this command.</p>

	Command	Purpose
Step 5	trust [cos dscp ip-precedence]	<p>Configure the trust state, which QoS uses to generate a CoS-based or DSCP-based QoS label.</p> <p>Note This command is mutually exclusive with the set command within the same policy map. If you enter the trust command, then skip Step 6.</p> <p>By default, the port is not trusted. If no keyword is specified when the command is entered, the default is dscp.</p> <p>The keywords have these meanings:</p> <ul style="list-style-type: none"> • cos—QoS derives the DSCP value by using the received or default port CoS value and the CoS-to-DSCP map. • dscp—QoS derives the DSCP value by using the DSCP value from the ingress packet. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. • ip-precedence—QoS derives the DSCP value by using the IP precedence value from the ingress packet and the IP-precedence-to-DSCP map. For non-IP packets that are tagged, QoS derives the DSCP value by using the received CoS value; for non-IP packets that are untagged, QoS derives the DSCP value by using the default port CoS value. In either case, the DSCP value is derived from the CoS-to-DSCP map. <p>For more information, see the “Configuring the CoS-to-DSCP Map” section on page 27-47.</p>
Step 6	set { ip dscp <i>new-dscp</i> ip precedence <i>new-precedence</i> }	<p>Classify IP traffic by setting a new value in the packet.</p> <ul style="list-style-type: none"> • For ip dscp <i>new-dscp</i>, enter a new DSCP value to be assigned to the classified traffic. The range is 0 to 63. • For ip precedence <i>new-precedence</i>, enter a new IP-precedence value to be assigned to the classified traffic. The range is 0 to 7.
Step 7	police <i>rate-bps burst-byte</i> [exceed-action { drop policed-dscp-transmit }]	<p>Define a policer for the classified traffic.</p> <p>By default, no policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 27-29.</p> <ul style="list-style-type: none"> • For <i>rate-bps</i>, specify average traffic rate in bits per second (bps). The range is 8000 to 1000000000. • For <i>burst-byte</i>, specify the normal burst size in bytes. The range is 8000 to 1000000. • (Optional) Specify the action to take when the rates are exceeded. Use the exceed-action drop keywords to drop the packet. Use the exceed-action policed-dscp-transmit keywords to mark down the DSCP value (by using the policed-DSCP map) and send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 27-49.

	Command	Purpose
Step 8	exit	Return to policy map configuration mode.
Step 9	exit	Return to global configuration mode.
Step 10	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to attach to the policy map. Valid interfaces include physical interfaces.
Step 11	service-policy input <i>policy-map-name</i>	Specify the policy-map name, and apply it to an ingress interface. Only one policy map per ingress interface is supported.
Step 12	end	Return to privileged EXEC mode.
Step 13	show policy-map [<i>policy-map-name</i> [class <i>class-map-name</i>]]	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an existing policy map, use the **no policy-map** *policy-map-name* global configuration command. To delete an existing class map, use the **no class** *class-map-name* policy-map configuration command. To return to the untrusted state, use the **no trust** policy-map configuration command. To remove an assigned DSCP or IP precedence value, use the **no set** {**ip dscp** *new-dscp* | **ip precedence** *new-precedence*} policy-map configuration command. To remove an existing policer, use the **no police** *rate-bps burst-byte* [**exceed-action** {**drop** | **policed-dscp-transmit**}] policy-map configuration command. To remove the policy map and interface association, use the **no service-policy input** *policy-map-name* interface configuration command.

This example shows how to create a policy map and attach it to an ingress interface. In the configuration, the IP standard ACL permits traffic from network 10.1.0.0. For traffic matching this classification, the DSCP value in the incoming packet is trusted. If the matched traffic exceeds an average traffic rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent:

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# policy-map flow1t
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police 48000 8000 exceed-action policed-dscp-transmit
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input flow1t
```

This example shows how to create a Layer 2 MAC ACL with two permit statements and attach it to an ingress interface. The first permit statement allows traffic from the host with MAC address 0001.0000.0001 destined for the host with MAC address 0002.0000.0001. The second permit statement allows only Ethertype XNS-IDP traffic from the host with MAC address 0001.0000.0002 destined for the host with MAC address 0002.0000.0002.

```
Switch(config)# mac access-list extended maclist1
Switch(config-ext-mac)# permit 0001.0000.0001 0.0.0 0002.0000.0001 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0002 0.0.0 0002.0000.0002 0.0.0 xns-idp
Switch(config-ext-mac)# exit
Switch(config)# mac access-list extended maclist2
Switch(config-ext-mac)# permit 0001.0000.0003 0.0.0 0002.0000.0003 0.0.0
Switch(config-ext-mac)# permit 0001.0000.0004 0.0.0 0002.0000.0004 0.0.0 aarp
```

```

Switch(config-ext-mac)# exit
Switch(config)# class-map macclass1
Switch(config-cmap)# match access-group maclist1
Switch(config-cmap)# exit
Switch(config)# policy-map macpolicy1
Switch(config-pmap)# class macclass1
Switch(config-pmap-c)# set ip dscp 63
Switch(config-pmap-c)# exit
Switch(config-pmap)# class macclass2 maclist2
Switch(config-pmap-c)# set ip dscp 45
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# mls qos trust cos
Switch(config-if)# service-policy input macpolicy1

```

Classifying, Policing, and Marking Traffic by Using Aggregate Policers

By using an aggregate policer, you can create a policer that is shared by multiple traffic classes within the same policy map. However, you cannot use the aggregate policer across different policy maps or interfaces.

Beginning in privileged EXEC mode, follow these steps to create an aggregate policer:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos aggregate-policer <i>aggregate-policer-name rate-bps burst-byte</i> exceed-action {drop policed-dscp-transmit}	<p>Define the policer parameters that can be applied to multiple traffic classes within the same policy map.</p> <p>By default, no aggregate policer is defined. For information on the number of policers supported, see the “Standard QoS Configuration Guidelines” section on page 27-29</p> <ul style="list-style-type: none"> For <i>aggregate-policer-name</i>, specify the name of the aggregate policer. For <i>rate-bps</i>, specify average traffic rate in bits per second (bps). The range is 8000 to 1000000000. For <i>burst-byte</i>, specify the normal burst size in bytes. The range is 8000 to 1000000. Specify the action to take when the rates are exceeded. Use the exceed-action drop keywords to drop the packet. Use the exceed-action policed-dscp-transmit keywords to mark down the DSCP value (by using the policed-DSCP map) and send the packet. For more information, see the “Configuring the Policed-DSCP Map” section on page 27-49.
Step 3	class-map [match-all match-any] <i>class-map-name</i>	Create a class map to classify traffic as necessary. For more information, see the “Classifying Traffic by Using Class Maps” section on page 27-40.
Step 4	policy-map policy-map-name	<p>Create a policy map by entering the policy map name, and enter policy-map configuration mode.</p> <p>For more information, see the “Classifying, Policing, and Marking Traffic by Using Policy Maps” section on page 27-42.</p>

	Command	Purpose
Step 5	class <i>class-map-name</i>	Define a traffic classification, and enter policy-map class configuration mode. For more information, see the “ Classifying, Policing, and Marking Traffic by Using Policy Maps ” section on page 27-42.
Step 6	police aggregate <i>aggregate-policer-name</i>	Apply an aggregate policer to multiple classes in the same policy map. For <i>aggregate-policer-name</i> , enter the name specified in Step 2.
Step 7	exit	Return to global configuration mode.
Step 8	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to attach to the policy map. Valid interfaces include physical interfaces.
Step 9	service-policy input <i>policy-map-name</i>	Specify the policy-map name, and apply it to an ingress interface. Only one policy map per ingress interface is supported.
Step 10	end	Return to privileged EXEC mode.
Step 11	show mls qos aggregate-policer [<i>aggregate-policer-name</i>]	Verify your entries.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the specified aggregate policer from a policy map, use the **no police aggregate** *aggregate-policer-name* policy map configuration mode. To delete an aggregate policer and its parameters, use the **no mls qos aggregate-policer** *aggregate-policer-name* global configuration command.

This example shows how to create an aggregate policer and attach it to multiple classes within a policy map. In the configuration, the IP ACLs permit traffic from network 10.1.0.0 and from host 11.3.1.1. For traffic coming from network 10.1.0.0, the DSCP in the incoming packets is trusted. For traffic coming from host 11.3.1.1, the DSCP in the packet is changed to 56. The traffic rate from the 10.1.0.0 network and from host 11.3.1.1 is policed. If the traffic exceeds an average rate of 48000 bps and a normal burst size of 8000 bytes, its DSCP is marked down (based on the policed-DSCP map) and sent. The policy map is attached to an ingress interface.

```
Switch(config)# access-list 1 permit 10.1.0.0 0.0.255.255
Switch(config)# access-list 2 permit 11.3.1.1
Switch(config)# mls qos aggregate-police transmit1 48000 8000 exceed-action
policed-dscp-transmit
Switch(config)# class-map ipclass1
Switch(config-cmap)# match access-group 1
Switch(config-cmap)# exit
Switch(config)# class-map ipclass2
Switch(config-cmap)# match access-group 2
Switch(config-cmap)# exit
Switch(config)# policy-map aggflow1
Switch(config-pmap)# class ipclass1
Switch(config-pmap-c)# trust dscp
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# class ipclass2
Switch(config-pmap-c)# set ip dscp 56
Switch(config-pmap-c)# police aggregate transmit1
Switch(config-pmap-c)# exit
Switch(config-pmap)# exit
```

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# service-policy input aggflow1
Switch(config-if)# exit
```

Configuring DSCP Maps

These sections describe how to configure the DSCP maps:

- [Configuring the CoS-to-DSCP Map, page 27-47](#) (optional)
- [Configuring the IP-Precedence-to-DSCP Map, page 27-48](#) (optional)
- [Configuring the Policed-DSCP Map, page 27-49](#) (optional, unless the null settings in the map are not appropriate)
- [Configuring the DSCP-to-CoS Map, page 27-50](#) (optional)
- [Configuring the DSCP-to-DSCP-Mutation Map, page 27-51](#) (optional, unless the null settings in the map are not appropriate)

All the maps, except the DSCP-to-DSCP-mutation map, are globally defined and are applied to all ports.

Configuring the CoS-to-DSCP Map

You use the CoS-to-DSCP map to map CoS values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

[Table 27-12](#) shows the default CoS-to-DSCP map.

Table 27-12 Default CoS-to-DSCP Map

CoS value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the CoS-to-DSCP map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map cos-dscp dscp1...dscp8	Modify the CoS-to-DSCP map. For <i>dscp1...dscp8</i> , enter eight DSCP values that correspond to CoS values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps cos-dscp	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos cos-dscp** global configuration command.

This example shows how to modify and display the CoS-to-DSCP map:

```
Switch(config)# mls qos map cos-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps cos-dscp

Cos-dscp map:
  cos:    0  1  2  3  4  5  6  7
-----
  dscp:   10 15 20 25 30 35 40 45
```

Configuring the IP-Precedence-to-DSCP Map

You use the IP-precedence-to-DSCP map to map IP precedence values in incoming packets to a DSCP value that QoS uses internally to represent the priority of the traffic.

Table 27-13 shows the default IP-precedence-to-DSCP map:

Table 27-13 Default IP-Precedence-to-DSCP Map

IP precedence value	0	1	2	3	4	5	6	7
DSCP value	0	8	16	24	32	40	48	56

If these values are not appropriate for your network, you need to modify them. Beginning in privileged EXEC mode, follow these steps to modify the IP-precedence-to-DSCP map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map ip-prec-dscp <i>dscp1...dscp8</i>	Modify the IP-precedence-to-DSCP map. For <i>dscp1...dscp8</i> , enter eight DSCP values that correspond to the IP precedence values 0 to 7. Separate each DSCP value with a space. The DSCP range is 0 to 63.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps ip-prec-dscp	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos ip-prec-dscp** global configuration command.

This example shows how to modify and display the IP-precedence-to-DSCP map:

```
Switch(config)# mls qos map ip-prec-dscp 10 15 20 25 30 35 40 45
Switch(config)# end
Switch# show mls qos maps ip-prec-dscp

IpPrecedence-dscp map:
  ipprec:  0  1  2  3  4  5  6  7
-----
  dscp:   10 15 20 25 30 35 40 45
```

Configuring the Policed-DSCP Map

You use the policed-DSCP map to mark down a DSCP value to a new value as the result of a policing and marking action.

The default policed-DSCP map is a null map, which maps an incoming DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the policed-DSCP map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map policed-dscp <i>dscp-list</i> to <i>mark-down-dscp</i>	Modify the policed-DSCP map. <ul style="list-style-type: none"> For <i>dscp-list</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. For <i>mark-down-dscp</i>, enter the corresponding policed (marked down) DSCP value.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps policed-dscp	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos policed-dscp** global configuration command.

This example shows how to map DSCP 50 to 57 to a marked-down DSCP value of 0:

```
Switch(config)# mls qos map policed-dscp 50 51 52 53 54 55 56 57 to 0
Switch(config)# end
Switch# show mls qos maps policed-dscp
Policed-dscp map:
  d1 :  d2 0  1  2  3  4  5  6  7  8  9
-----
  0 :    00 01 02 03 04 05 06 07 08 09
  1 :    10 11 12 13 14 15 16 17 18 19
  2 :    20 21 22 23 24 25 26 27 28 29
  3 :    30 31 32 33 34 35 36 37 38 39
  4 :    40 41 42 43 44 45 46 47 48 49
  5 :    00 00 00 00 00 00 00 00 58 59
  6 :    60 61 62 63
```



Note

In this policed-DSCP map, the marked-down DSCP values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the marked-down value. For example, an original DSCP value of 53 corresponds to a marked-down DSCP value of 0.

Configuring the DSCP-to-CoS Map

You use the DSCP-to-CoS map to generate a CoS value, which is used to select one of the four egress queues.

Table 27-14 shows the default DSCP-to-CoS map.

Table 27-14 Default DSCP-to-CoS Map

DSCP value	0–7	8–15	16–23	24–31	32–39	40–47	48–55	56–63
CoS value	0	1	2	3	4	5	6	7

If these values are not appropriate for your network, you need to modify them.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-CoS map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map dscp-cos <i>dscp-list</i> to <i>cos</i>	Modify the DSCP-to-CoS map. <ul style="list-style-type: none"> For <i>dscp-list</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. For <i>cos</i>, enter the CoS value to which the DSCP values correspond. The DSCP range is 0 to 63; the CoS range is 0 to 7.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps dscp-to-cos	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos dscp-cos** global configuration command.

This example shows how to map DSCP values 0, 8, 16, 24, 32, 40, 48, and 50 to CoS value 0 and to display the map:

```
Switch(config)# mls qos map dscp-cos 0 8 16 24 32 40 48 50 to 0
Switch(config)# end
Switch# show mls qos maps dscp-cos
Dscp-cos map:
  d1 : d2 0  1  2  3  4  5  6  7  8  9
-----
  0 :   00 00 00 00 00 00 00 00 00 01
  1 :   01 01 01 01 01 01 00 02 02 02
  2 :   02 02 02 02 00 03 03 03 03 03
  3 :   03 03 00 04 04 04 04 04 04 04
  4 :   00 05 05 05 05 05 05 05 00 06
  5 :   00 06 06 06 06 06 07 07 07 07
  6 :   07 07 07 07
```



Note

In the above DSCP-to-CoS map, the CoS values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the DSCP; the d2 row specifies the least-significant digit of the DSCP. The intersection of the d1 and d2 values provides the CoS value. For example, in the DSCP-to-CoS map, a DSCP value of 08 corresponds to a CoS value of 0.

Configuring the DSCP-to-DSCP-Mutation Map

If two QoS domains have different DSCP definitions, use the DSCP-to-DSCP-mutation map to translate one set of DSCP values to match the definition of another domain. You apply the DSCP-to-DSCP-mutation map to the receiving interface (ingress mutation) at the boundary of a QoS administrative domain.

With ingress mutation, the new DSCP value overwrites the one in the packet, and QoS treats the packet with this new value. The switch sends the packet out the interface with the new DSCP value.

You can configure multiple DSCP-to-DSCP-mutation maps on an ingress port. The default DSCP-to-DSCP-mutation map is a null map, which maps an incoming DSCP value to the same DSCP value.

Beginning in privileged EXEC mode, follow these steps to modify the DSCP-to-DSCP-mutation map. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos map dscp-mutation <i>dscp-mutation-name in-dscp to out-dscp</i>	Modify the DSCP-to-DSCP-mutation map. <ul style="list-style-type: none"> For <i>dscp-mutation-name</i>, enter the mutation map name. You can create more than one map by specifying a new name. For <i>in-dscp</i>, enter up to eight DSCP values separated by spaces. Then enter the to keyword. For <i>out-dscp</i>, enter a single DSCP value. The DSCP range is 0 to 63.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to which to attach the map. Valid interfaces include physical interfaces.
Step 4	mls qos trust dscp	Configure the ingress port as a DSCP-trusted port. By default, the port is not trusted.
Step 5	mls qos dscp-mutation <i>dscp-mutation-name</i>	Apply the map to the specified ingress DSCP-trusted port. For <i>dscp-mutation-name</i> , enter the mutation map name specified in Step 2.
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos maps dscp-mutation	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default map, use the **no mls qos dscp-mutation** *dscp-mutation-name* global configuration command.

This example shows how to define the DSCP-to-DSCP-mutation map. All the entries that are not explicitly configured are not modified (remains as specified in the null map):

```
Switch(config)# mls qos map dscp-mutation mutation1 1 2 3 4 5 6 7 to 0
Switch(config)# mls qos map dscp-mutation mutation1 8 9 10 11 12 13 to 10
Switch(config)# mls qos map dscp-mutation mutation1 20 21 22 to 20
Switch(config)# mls qos map dscp-mutation mutation1 30 31 32 33 34 to 30
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# mls qos trust dscp
Switch(config-if)# mls qos dscp-mutation mutation1
Switch(config-if)# end
Switch# show mls qos maps dscp-mutation mutation1
Dscp-dscp mutation map:
mutation1:
d1 : d2 0 1 2 3 4 5 6 7 8 9
-----
0 : 00 00 00 00 00 00 00 00 10 10
1 : 10 10 10 10 14 15 16 17 18 19
2 : 20 20 20 23 24 25 26 27 28 29
3 : 30 30 30 30 30 35 36 37 38 39
4 : 40 41 42 43 44 45 46 47 48 49
5 : 50 51 52 53 54 55 56 57 58 59
6 : 60 61 62 63
```



Note

In the above DSCP-to-DSCP-mutation map, the mutated values are shown in the body of the matrix. The d1 column specifies the most-significant digit of the original DSCP; the d2 row specifies the least-significant digit of the original DSCP. The intersection of the d1 and d2 values provides the mutated value. For example, a DSCP value of 12 corresponds to a mutated value of 10.

Configuring Ingress Queue Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are assigned (by DSCP or CoS value) to each queue?
- What drop percentage thresholds apply to each queue, and which CoS or DSCP values map to each threshold?
- How much of the available buffer space is allocated between the queues?
- How much of the available bandwidth is allocated between the queues?
- Is there traffic (such as voice) that should be given high priority?

These sections describe how to configure ingress queue characteristics:

- [Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds, page 27-53](#) (optional)
- [Allocating Buffer Space Between the Ingress Queues, page 27-54](#) (optional)
- [Allocating Bandwidth Between the Ingress Queues, page 27-55](#) (optional)
- [Configuring the Ingress Priority Queue, page 27-56](#) (optional)

Mapping DSCP or CoS Values to an Ingress Queue and Setting WTD Thresholds

You can prioritize traffic by placing packets with particular DSCPs or CoSs into certain queues and adjusting the queue thresholds so that packets with lower priorities are dropped.

Beginning in privileged EXEC mode, follow these steps to map DSCP or CoS values to an ingress queue and to set WTD thresholds. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input dscp-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>dscp1...dscp8</i> or mls qos srr-queue input cos-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>cos1...cos8</i>	Map DSCP or CoS values to an ingress queue and to a threshold ID. By default, DSCP values 0–39 and 48–63 are mapped to queue 1 and threshold 1. DSCP values 40–47 are mapped to queue 2 and threshold 1. By default, CoS values 0–4, 6, and 7 are mapped to queue 1 and threshold 1. CoS value 5 is mapped to queue 2 and threshold 1. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 2. For <i>threshold-id</i>, the range is 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. For <i>dscp1...dscp8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 63. For <i>cos1...cos8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 7.
Step 3	mls qos srr-queue input threshold <i>queue-id</i> <i>threshold-percentage1</i> <i>threshold-percentage2</i>	Assign the two WTD threshold percentages for (threshold 1 and 2) to an ingress queue. The default, both thresholds are set to 100 percent. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 2. For <i>threshold-percentage1 threshold-percentage2</i>, the range is 1 to 100. Separate each value with a space. Each threshold value is a percentage of the total number of queue descriptors allocated for the queue.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos maps	Verify your entries. The DSCP input queue threshold map appears as a matrix. The d1 column specifies the most-significant digit of the DSCP number; the d2 row specifies the least-significant digit in the DSCP number. The intersection of the d1 and the d2 values provides the queue ID and threshold ID; for example, queue 2 and threshold 1 (02-01). The CoS input queue threshold map shows the CoS value in the top row and the corresponding queue ID and threshold ID in the second row; for example, queue 2 and threshold 2 (2-2).
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default CoS input queue threshold map or the default DSCP input queue threshold map, use the **no mls qos srr-queue input cos-map** or the **no mls qos srr-queue input dscp-map** global configuration command. To return to the default WTD threshold percentages, use the **no mls qos srr-queue input threshold** *queue-id* global configuration command.

This example shows how to map DSCP values 0 to 6 to ingress queue 1 and to threshold 1 with a drop threshold of 50 percent. It maps DSCP values 20 to 26 to ingress queue 1 and to threshold 2 with a drop threshold of 70 percent:

```
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 1 0 1 2 3 4 5 6
Switch(config)# mls qos srr-queue input dscp-map queue 1 threshold 2 20 21 22 23 24 25 26
Switch(config)# mls qos srr-queue input threshold 1 50 70
```

In this example, the DSCP values (0 to 6) are assigned the WTD threshold of 50 percent and will be dropped sooner than the DSCP values (20 to 26) assigned to the WTD threshold of 70 percent.

Allocating Buffer Space Between the Ingress Queues

You define the ratio (allocate the amount of space) with which to divide the ingress buffers between the two queues. The buffer and the bandwidth allocation determine how much data can be buffered before packets are dropped.

Beginning in privileged EXEC mode, follow these steps to allocate the buffers between the ingress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input buffers <i>percentage1 percentage2</i>	Allocate the buffers between the ingress queues By default 90 percent of the buffers are allocated to queue 1, and 10 percent of the buffers are allocated to queue 2. For <i>percentage1 percentage2</i> , the range is 0 to 100. Separate each value with a space. You should allocate the buffers so that the queues can handle any incoming bursty traffic.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface buffer or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input buffers** global configuration command.

This example shows how to allocate 60 percent of the buffer space to ingress queue 1 and 40 percent of the buffer space to ingress queue 2:

```
Switch(config)# mls qos srr-queue input buffers 60 40
```

Allocating Bandwidth Between the Ingress Queues

You need to specify how much of the available bandwidth is allocated between the ingress queues. The ratio of the weights is the ratio of the frequency in which the SRR scheduler sends packets from each queue. The bandwidth and the buffer allocation determine how much data can be buffered before packets are dropped. On ingress queues, SRR operates only in shared mode.

Beginning in privileged EXEC mode, follow these steps to allocate bandwidth between the ingress queues. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input bandwidth <i>weight1 weight2</i>	Assign shared round robin weights to the ingress queues. The default setting for <i>weight1</i> and <i>weight2</i> is 4 (1/2 of the bandwidth is equally shared between the two queues). For <i>weight1</i> and <i>weight2</i> , the range is 1 to 100. Separate each value with a space. SRR services the priority queue for its configured weight as specified by the bandwidth keyword in the mls qos srr-queue input priority-queue queue-id bandwidth weight global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the mls qos srr-queue input bandwidth weight1 weight2 global configuration command. For more information, see the “Configuring the Ingress Priority Queue” section on page 27-56 .
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface queueing or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input bandwidth** global configuration command.

This example shows how to assign the ingress bandwidth to the queues. Priority queueing is disabled, and the shared bandwidth ratio allocated to queue 1 is 25/(25+75) and to queue 2 is 75/(25+75):

```
Switch(config)# mls qos srr-queue input priority-queue 2 bandwidth 0
Switch(config)# mls qos srr-queue input bandwidth 25 75
```

Configuring the Ingress Priority Queue

You should use the priority queue only for traffic that needs to be expedited (for example, voice traffic, which needs minimum delay and jitter).

The priority queue is guaranteed part of the bandwidth to reduce the delay and jitter under heavy network traffic on an oversubscribed ring (when there is more traffic than the backplane can carry, and the queues are full and dropping frames).

SRR services the priority queue for its configured weight as specified by the **bandwidth** keyword in the **mls qos srr-queue input priority-queue *queue-id* bandwidth *weight*** global configuration command. Then, SRR shares the remaining bandwidth with both ingress queues and services them as specified by the weights configured with the **mls qos srr-queue input bandwidth *weight1 weight2*** global configuration command.

Beginning in privileged EXEC mode, follow these steps to configure the priority queue. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue input priority-queue <i>queue-id</i> bandwidth <i>weight</i>	Assign a queue as the priority queue and guarantee bandwidth on the internal ring if the ring is congested. By default, the priority queue is queue 2, and 10 percent of the bandwidth is allocated to it. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 2. For bandwidth <i>weight</i>, assign the bandwidth percentage of the internal ring. The range is 0 to 40. The amount of bandwidth that can be guaranteed is restricted because a large value affects the entire ring and can degrade performance.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos interface queueing or show mls qos input-queue	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos srr-queue input priority-queue *queue-id*** global configuration command. To disable priority queueing, set the bandwidth weight to 0, for example, **mls qos srr-queue input priority-queue *queue-id* bandwidth 0**.

This example shows how to assign the ingress bandwidths to the queues. Queue 1 is the priority queue with 10 percent of the bandwidth allocated to it. The bandwidth ratios allocated to queues 1 and 2 is 4/(4+4). SRR services queue 1 (the priority queue) first for its configured 10 percent bandwidth. Then SRR equally shares the remaining 90 percent of the bandwidth between queues 1 and 2 by allocating 45 percent to each queue:

```
Switch(config)# mls qos srr-queue input priority-queue 1 bandwidth 10
Switch(config)# mls qos srr-queue input bandwidth 4 4
```

Configuring Egress Queue Characteristics

Depending on the complexity of your network and your QoS solution, you might need to perform all of the tasks in the next sections. You will need to make decisions about these characteristics:

- Which packets are mapped by DSCP or CoS value to each queue and threshold ID?
- What drop percentage thresholds apply to the queue-set (four egress queues per port), and how much reserved and maximum memory is needed for the traffic type?
- How much of the fixed buffer space is allocated to the queue-set?
- Does the bandwidth of the port need to be rate limited?
- How often should the egress queues be serviced and which technique (shaped, shared, or both) should be used?

These sections describe how to configure egress queue characteristics:

- [Configuration Guidelines, page 27-57](#)
- [Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set, page 27-58](#) (optional)
- [Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID, page 27-60](#) (optional)
- [Configuring SRR Shaped Weights on Egress Queues, page 27-61](#) (optional)
- [Configuring SRR Shared Weights on Egress Queues, page 27-62](#) (optional)
- [Configuring the Egress Expedite Queue, page 27-63](#) (optional)
- [Limiting the Bandwidth on an Egress Interface, page 27-64](#) (optional)

Configuration Guidelines

Follow these guidelines when the expedite queue is enabled or the egress queues are serviced based on their SRR weights:

- If the egress expedite queue is enabled, it overrides the SRR shaped and shared weights for queue 1.
- If the egress expedite queue is disabled and the SRR shaped and shared weights are configured, the shaped mode overrides the shared mode for queue 1, and SRR services this queue in shaped mode.
- If the egress expedite queue is disabled and the SRR shaped weights are not configured, SRR services this queue in shared mode.

Allocating Buffer Space to and Setting WTD Thresholds for an Egress Queue-Set

You can guarantee the availability of buffers, set WTD thresholds, and configure the maximum memory allocation for a queue-set by using the **mls qos queue-set output *qset-id* threshold *queue-id* drop-threshold1 drop-threshold2 reserved-threshold maximum-threshold** global configuration command.

Each threshold value is a percentage of the queue's allocated memory, which you specify by using the **mls qos queue-set output *qset-id* buffers *allocation1* ... *allocation4*** global configuration command. The queues use WTD to support distinct drop percentages for different traffic classes.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to configure the memory allocation and drop thresholds for a queue-set. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos queue-set output <i>qset-id</i> buffers <i>allocation1</i> ... <i>allocation4</i>	<p>Allocate buffers to a queue-set.</p> <p>By default, all allocation values are equally mapped among the four queues (25, 25, 25, 25). Each queue has 1/4 of the buffer space.</p> <ul style="list-style-type: none"> For <i>qset-id</i>, enter the ID of the queue-set. The range is 1 to 2. Each port belongs to a queue-set, which defines all the characteristics of the four egress queues per port. For <i>allocation1</i> ... <i>allocation4</i>, specify four percentages, one for each queue in the queue-set. The range is 0 to 100. Separate each value with a space. <p>Allocate buffers according to the importance of the traffic; for example, give a large percentage of the buffer to the queue with the highest-priority traffic.</p>

	Command	Purpose
Step 3	mls qos queue-set output <i>qset-id</i> threshold <i>queue-id drop-threshold1</i> <i>drop-threshold2 reserved-threshold</i> <i>maximum-threshold</i>	<p>Configure the WTD thresholds, guarantee the availability of buffers, and configure the maximum memory allocation for the queue-set (four egress queues per port).</p> <p>By default, the WTD thresholds for queues 1, 3, and 4 are set to 100 percent. The thresholds for queue 2 are set to 50 percent. The reserved thresholds for queues 1, 3, and 4 are set to 50 percent. The reserved threshold for queue 2 is set to 100 percent. The maximum thresholds for all queues are set to 400 percent.</p> <ul style="list-style-type: none"> For <i>qset-id</i>, enter the ID of the queue-set specified in Step 2. The range is 1 to 2. For <i>queue-id</i>, enter the specific queue in the queue-set on which the command is performed. The range is 1 to 4. For <i>drop-threshold1 drop-threshold2</i>, specify the two WTD thresholds expressed as a percentage of the queue's allocated memory. The range is 1 to 400 percent. For <i>reserved-threshold</i>, enter the amount of memory to be guaranteed (reserved) for the queue expressed as a percentage of the allocated memory. The range is 1 to 100 percent. For <i>maximum-threshold</i>, enable a queue in the full condition to obtain more buffers than are reserved for it. This is the maximum memory the queue can have before the packets are dropped if the common pool is not empty. The range is 1 to 400 percent.
Step 4	interface <i>interface-id</i>	Specify the interface of the outbound traffic, and enter interface configuration mode.
Step 5	queue-set <i>qset-id</i>	<p>Map the port to a queue-set.</p> <p>For <i>qset-id</i>, enter the ID of the queue-set specified in Step 2. The range is 1 to 2. The default is 1.</p>
Step 6	end	Return to privileged EXEC mode.
Step 7	show mls qos interface [<i>interface-id</i>] buffers	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no mls qos queue-set output** *qset-id buffers* global configuration command. To return to the default WTD threshold percentages, use the **no mls qos queue-set output** *qset-id threshold* [*queue-id*] global configuration command.

This example shows how to map Gigabit Ethernet interface 0/1 to queue-set 2. It allocates 40 percent of the buffer space to egress queue 1 and 20 percent to egress queues 2, 3, and 4. It configures the drop thresholds for queue 2 to 40 and 60 percent of the allocated memory, guarantees (reserves) 100 percent of the allocated memory, and configures 200 percent as the maximum memory this queue can have before packets are dropped:

```
Switch(config)# mls qos queue-set output 2 buffers 40 20 20 20
Switch(config)# mls qos queue-set output 2 threshold 2 40 60 100 200
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# queue-set 2
```

Mapping DSCP or CoS Values to an Egress Queue and to a Threshold ID

You can prioritize traffic by placing packets with particular DSCPs or costs of service into certain queues and adjusting the queue thresholds so that packets with lower priorities are dropped.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to map DSCP or CoS values to an egress queue and to a threshold ID. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos srr-queue output dscp-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>dscp1...dscp8</i> or mls qos srr-queue output cos-map queue <i>queue-id</i> threshold <i>threshold-id</i> <i>cos1...cos8</i>	Map DSCP or CoS values to an egress queue and to a threshold ID. By default, DSCP values 0–15 are mapped to queue 2 and threshold 1. DSCP values 16–31 are mapped to queue 3 and threshold 1. DSCP values 32–39 and 48–63 are mapped to queue 4 and threshold 1. DSCP values 40–47 are mapped to queue 1 and threshold 1. By default, CoS values 0 and 1 are mapped to queue 2 and threshold 1. CoS values 2 and 3 are mapped to queue 3 and threshold 1. CoS values 4, 6, and 7 are mapped to queue 4 and threshold 1. CoS value 5 is mapped to queue 1 and threshold 1. <ul style="list-style-type: none"> For <i>queue-id</i>, the range is 1 to 4. For <i>threshold-id</i>, the range is 1 to 3. The drop-threshold percentage for threshold 3 is predefined. It is set to the queue-full state. For <i>dscp1...dscp8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 63. For <i>cos1...cos8</i>, enter up to eight values, and separate each value with a space. The range is 0 to 7.
Step 3	end	Return to privileged EXEC mode.
Step 4	show mls qos maps	Verify your entries. The DSCP output queue threshold map appears as a matrix. The d1 column specifies the most-significant digit of the DSCP number; the d2 row specifies the least-significant digit in the DSCP number. The intersection of the d1 and the d2 values provides the queue ID and threshold ID; for example, queue 2 and threshold 1 (02-01). The CoS output queue threshold map shows the CoS value in the top row and the corresponding queue ID and threshold ID in the second row; for example, queue 2 and threshold 2 (2-2).
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default DSCP output queue threshold map or the default CoS output queue threshold map, use the **no mls qos srr-queue output dscp-map** or the **no mls qos srr-queue output cos-map** global configuration command.

This example shows how to map DSCP values 10 and 11 to egress queue 1 and to threshold 2:

```
Switch(config)# mls qos srr-queue output dscp-map queue 1 threshold 2 10 11
```

Configuring SRR Shaped Weights on Egress Queues

You can specify how much of the available bandwidth is allocated to each queue. The ratio of the weights is the ratio of frequency in which the SRR scheduler sends packets from each queue.

You can configure the egress queues for shaped or shared weights, or both. Use shaping to smooth bursty traffic or to provide a smoother output over time. For information about, see the [“SRR Shaping and Sharing”](#) section on page 27-12. For information about shared weights, see the [“Configuring SRR Shared Weights on Egress Queues”](#) section on page 27-62.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to assign the shaped weights and to enable bandwidth shaping on the four egress queues mapped to a port. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface of the outbound traffic, and enter interface configuration mode.
Step 3	srr-queue bandwidth shape <i>weight1 weight2 weight3 weight4</i>	<p>Assign SRR weights to the egress queues.</p> <p>By default, <i>weight1</i> is set to 25; <i>weight2</i>, <i>weight3</i>, and <i>weight4</i> are set to 0, and these queues are in shared mode.</p> <p>For <i>weight1 weight2 weight3 weight4</i>, enter the weights to determine the percentage of the port that is shaped. The inverse ratio ($1/\text{weight}$) determines the shaping bandwidth for this queue. Separate each value with a space. The range is 0 to 65535.</p> <p>If you configure a weight of 0, the corresponding queue operates in shared mode. The weight specified with the srr-queue bandwidth shape command is ignored, and the weights specified with the srr-queue bandwidth share interface configuration command for a queue come into effect. When configuring queues in the same queue-set for both shaping and sharing, make sure that you configure the lowest number queue for shaping.</p> <p>The shaped mode overrides the shared mode.</p>
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface <i>interface-id</i> queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth shape** interface configuration command.

This example shows how to configure bandwidth shaping on queue 1. Because the weight ratios for queues 2, 3, and 4 are set to 0, these queues operate in shared mode. The bandwidth weight for queue 1 is 1/8, which is 12.5 percent:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth shape 8 0 0 0
```

Configuring SRR Shared Weights on Egress Queues

In shared mode, the queues share the bandwidth among them according to the configured weights. The bandwidth is guaranteed at this level but not limited to it. For example, if a queue empties and does not require a share of the link, the remaining queues can expand into the unused bandwidth and share it among them. With sharing, the ratio of the weights determines the frequency of dequeuing; the absolute values are meaningless.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to assign the shared weights and to enable bandwidth sharing on the four egress queues mapped to a port. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface of the outbound traffic, and enter interface configuration mode.
Step 3	srr-queue bandwidth share <i>weight1 weight2 weight3 weight4</i>	Assign SRR weights to the egress queues. By default, all four weights are 25 (1/4 of the bandwidth is allocated to each queue). For <i>weight1 weight2 weight3 weight4</i> , enter the weights to determine the ratio of the frequency in which the SRR scheduler sends packets. Separate each value with a space. The range is 1 to 255.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface <i>interface-id</i> queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth share** interface configuration command.

This example shows how to configure the weight ratio of the SRR scheduler running on egress port Gigabit Ethernet 0/1. Four queues are used, and the bandwidth ratio allocated for each queue in shared mode is 1/(1+2+3+4), 2/(1+2+3+4), 3/(1+2+3+4), and 4/(1+2+3+4), which is 10 percent, 20 percent, 30 percent, and 40 percent for queues 1, 2, 3, and 4. This means that queue 4 has four times the bandwidth of queue 1, twice the bandwidth of queue 2, and one-and-a-third times the bandwidth of queue 3.

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth share 1 2 3 4
```

Configuring the Egress Expedite Queue

Beginning in Cisco IOS Release 12.1(19)EA1, you can ensure that certain packets have priority over all others by queuing them in the egress expedite queue. SRR services this queue until it is empty before servicing the other queues.

Beginning in privileged EXEC mode, follow these steps to enable the egress expedite queue. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	mls qos	Enable QoS on a switch.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the egress interface.
Step 4	priority-queue out	Enable the egress expedite queue, which is disabled by default. When you configure this command, the SRR weight and queue size ratios are affected because there is one fewer queue participating in SRR. This means that <i>weight1</i> in the srr-queue bandwidth shape or the srr-queue bandwidth share command is ignored (not used in the ratio calculation).
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable the egress expedite queue, use the **no priority-queue out** interface configuration command.

This example shows how to enable the egress expedite queue when the SRR weights are configured. The egress expedite queue overrides the configured SRR weights.

```
Switch(config)# interface gigabitethernet0/4
Switch(config-if)# srr-queue bandwidth shape 25 0 0 0
Switch(config-if)# srr-queue bandwidth share 30 20 25 25
Switch(config-if)# priority-queue out
Switch(config-if)# end
```

Limiting the Bandwidth on an Egress Interface

You can limit the bandwidth on an egress interface. For example, if a customer pays only for a small percentage of a high-speed link, you can limit the bandwidth to that amount.



Note

The egress queue default settings are suitable for most situations. You should change them only when you have a thorough understanding of the egress queues and if these settings do not meet your QoS solution.

Beginning in privileged EXEC mode, follow these steps to limit the bandwidth on an egress interface. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be rate limited, and enter interface configuration mode.
Step 3	srr-queue bandwidth limit <i>weight1</i>	Specify the percentage of the port speed to which the port should be limited. The range is 10 to 90. By default, the port is not rate limited and is set to 100 percent.
Step 4	end	Return to privileged EXEC mode.
Step 5	show mls qos interface [<i>interface-id</i>] queueing	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default setting, use the **no srr-queue bandwidth limit** interface configuration command.

This example shows how to limit the bandwidth on Gigabit Ethernet interface 0/1 to 80 percent:

```
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# srr-queue bandwidth limit 80
```

When you configure this command to 80 percent, the port is idle 20 percent of the time. The line rate drops to 80 percent of the connected speed, which is 800 Mbps. These values are not exact because the hardware adjusts the line rate in increments of six.

Displaying Standard QoS Information

To display standard QoS information, use one or more of the privileged EXEC commands in [Table 27-15](#):

Table 27-15 Commands for Displaying Standard QoS Information

Command	Purpose
show class-map [<i>class-map-name</i>]	Display QoS class maps, which define the match criteria to classify traffic.
show mls qos	Display global QoS configuration information.
show mls qos aggregate-policer [<i>aggregate-policer-name</i>]	Display the aggregate policer configuration.
show mls qos input-queue	Display QoS settings for the ingress queues.
show mls qos interface [<i>interface-id</i>] [buffers policers queueing statistics]	Display QoS information at the interface level, including the buffer allocation, which interfaces have configured policers, the queueing strategy, and the ingress and egress statistics.
show mls qos maps [cos-dscp cos-input-q cos-output-q dscp-cos dscp-input-q dscp-mutation <i>dscp-mutation-name</i> dscp-output-q ip-prec-dscp policed-dscp]	Display QoS mapping information.
show mls qos queue-set [<i>qset-id</i>]	Display QoS settings for the egress queues.
show policy-map [<i>policy-map-name</i>] [class <i>class-map-name</i>]]	Display QoS policy maps, which define classification criteria for incoming traffic. Note Do not use the show policy-map interface privileged EXEC command to display classification information for incoming traffic. The interface keyword is not supported, and the statistics shown in the display should be ignored.



Configuring EtherChannels

This chapter describes how to configure EtherChannels on Layer 2 interfaces on the Catalyst 2970 switch. EtherChannel provides fault-tolerant high-speed links between switches, routers, and servers. You can use it to increase the bandwidth between the wiring closets and the data center, and you can deploy it anywhere in the network where bottlenecks are likely to occur. EtherChannel provides automatic recovery for the loss of a link by redistributing the load across the remaining links. If a link fails, EtherChannel redirects traffic from the failed link to the remaining links in the channel without intervention.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release.

This chapter consists of these sections:

- [Understanding EtherChannels, page 28-1](#)
- [Configuring EtherChannels, page 28-8](#)
- [Displaying EtherChannel, PAgP, and LACP Status, page 28-17](#)

Understanding EtherChannels

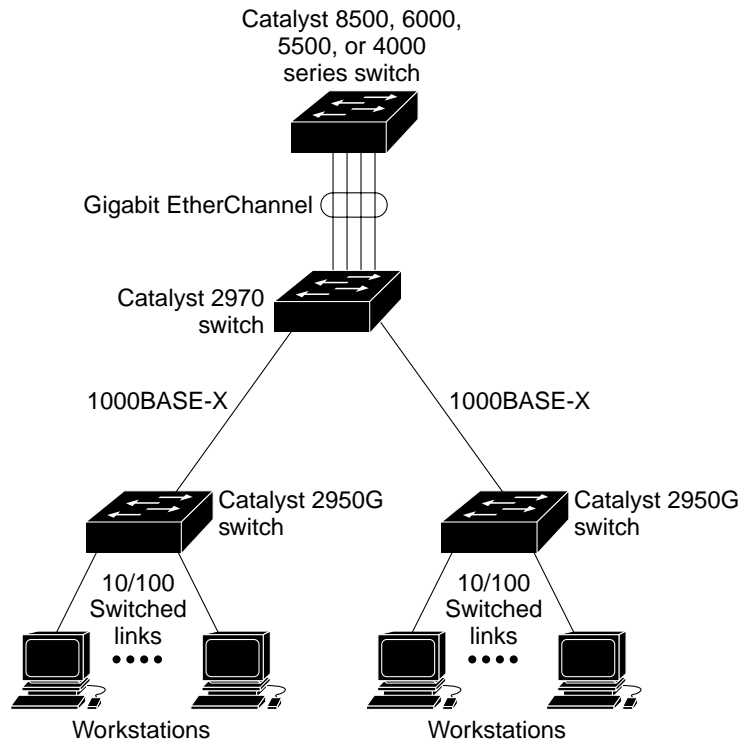
These sections describe how EtherChannels work:

- [EtherChannel Overview, page 28-2](#)
- [Port-Channel Interfaces, page 28-3](#)
- [Port Aggregation Protocol, page 28-3](#)
- [Link Aggregation Control Protocol, page 28-5](#)
- [Load Balancing and Forwarding Methods, page 28-6](#)

EtherChannel Overview

An EtherChannel consists of individual Gigabit Ethernet links bundled into a single logical link as shown in [Figure 28-1](#).

Figure 28-1 Typical EtherChannel Configuration



The EtherChannel provides full-duplex bandwidth up to 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host.

Each EtherChannel can consist of up to eight compatibly configured Ethernet interfaces. All interfaces in each EtherChannel must be configured as Layer 2 interfaces. For Catalyst 2970 switches, the number of EtherChannels is limited to 12. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 28-9.

If a link within an EtherChannel fails, traffic previously carried over that failed link changes to the remaining links within the EtherChannel. A trap is sent for a failure, identifying the switch, the EtherChannel, and the failed link. Inbound broadcast and multicast packets on one link in an EtherChannel are blocked from returning on any other link of the EtherChannel.

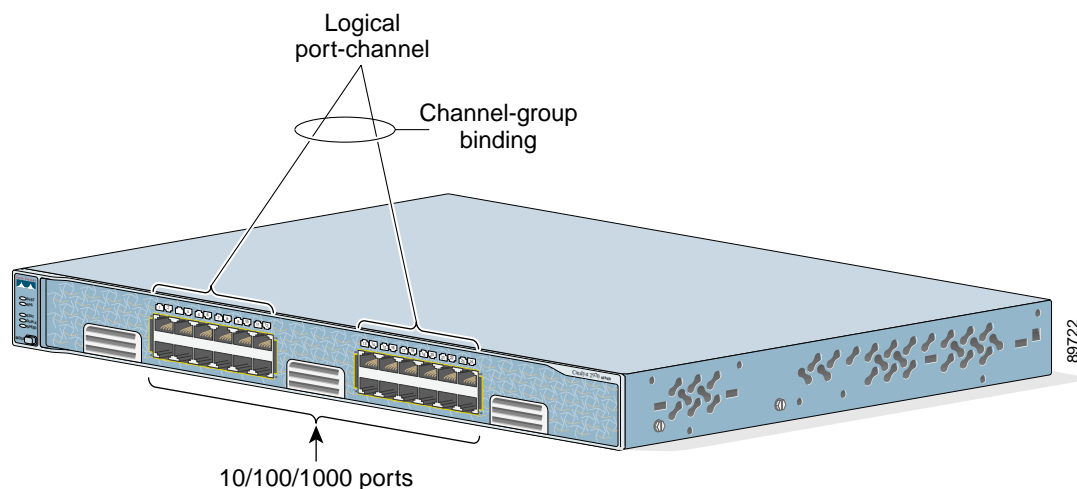
Port-Channel Interfaces

When you create a Layer 2 EtherChannel, a port-channel logical interface is involved. You can create the EtherChannel in these ways:

- Use the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface when the channel group gets its first physical interface. The **channel-group** command binds the physical (10/100/1000 ports) and the logical ports together as shown in Figure 28-2.
- Use the **interface port-channel** *port-channel-number* global configuration command to manually create the port-channel logical interface. Then use the **channel-group** *channel-group-number* interface configuration command to bind the logical interface to a physical port. The *channel-group-number* can be the same as the *port-channel-number*, or you can use a new number. If you use a new number, the **channel-group** command dynamically creates a new port channel.

Each EtherChannel has a port-channel logical interface numbered from 1 to 12. This port-channel interface number corresponds to the one specified with the **channel-group** interface configuration command.

Figure 28-2 Relationship of Physical Ports, Logical Port Channels, and Channel Groups



After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface. Configuration changes applied to the physical interface affect only the interface where you apply the configuration. To change the parameters of all ports in an EtherChannel, apply configuration commands to the port-channel interface, for example, spanning-tree commands or commands to configure a Layer 2 EtherChannel as a trunk.

Port Aggregation Protocol

The Port Aggregation Protocol (PAgP) is a Cisco-proprietary protocol that can be run only on Cisco switches and on those switches licensed by vendors to support PAgP. PAgP facilitates the automatic creation of EtherChannels by exchanging PAgP packets between Ethernet interfaces.

By using PAgP, the switch learns the identity of partners capable of supporting PAgP and the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port). Similarly configured interfaces are grouped based on hardware,

administrative, and port parameter constraints. For example, PAgP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, PAgP adds the group to the spanning tree as a single switch port.

PAgP Modes

Table 28-1 shows the user-configurable EtherChannel PAgP modes for the **channel-group** interface configuration command.

Table 28-1 EtherChannel PAgP Modes

Mode	Description
auto	Places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. This setting minimizes the transmission of PAgP packets.
desirable	Places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets.
on	Forces the interface to channel without PAgP (or the Link Aggregation Control Protocol [LACP]). With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode.

Switch interfaces exchange PAgP packets only with partner interfaces configured in the **auto** or **desirable** modes. Interfaces configured in the **on** mode do not exchange PAgP packets.

Both the **auto** and **desirable** modes allow interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different PAgP modes as long as the modes are compatible. For example:

- An interface in the **desirable** mode can form an EtherChannel with another interface that is in the **desirable** or **auto** mode.
- An interface in the **auto** mode can form an EtherChannel with another interface in the **desirable** mode.

An interface in the **auto** mode cannot form an EtherChannel with another interface that is also in the **auto** mode because neither interface starts PAgP negotiation.

An interface in the **on** mode that is added to a port channel is forced to have the same characteristics as the already existing **on** mode interfaces in the channel.



Caution

You should exercise care when setting the mode to **on** (manual configuration). All ports configured in the **on** mode are bundled in the same group and are forced to have similar characteristics. If the group is misconfigured, packet loss or spanning-tree loops might occur.

If your switch is connected to a partner that is PAgP-capable, you can configure the switch interface for nonsilent operation by using the **non-silent** keyword. If you do not specify **non-silent** with the **auto** or **desirable** mode, silent mode is assumed.

Use the silent mode when the switch is connected to a device that is not PAgP-capable and seldom, if ever, sends packets. An example of a silent partner is a file server or a packet analyzer that is not generating traffic. In this case, running PAgP on a physical port connected to a silent partner prevents that switch port from ever becoming operational. However, the silent setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

PAgP Interaction with Other Features

The Dynamic Trunking Protocol (DTP) and the Cisco Discovery Protocol (CDP) send and receive packets over the physical interfaces in the EtherChannel. Trunk ports send and receive PAgP protocol data units (PDUs) on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

PAgP sends and receives PAgP PDUs only from interfaces that are up and have PAgP enabled for the auto or desirable mode.

Link Aggregation Control Protocol

The LACP is defined in IEEE 802.3AD and enables Cisco switches to manage Ethernet channels between switches that conform to the 802.3AD protocol. LACP facilitates the automatic creation of EtherChannels by exchanging LACP packets between Ethernet interfaces.

By using LACP, the switch learns the identity of partners capable of supporting LACP and the capabilities of each interface. It then dynamically groups similarly configured interfaces into a single logical link (channel or aggregate port). Similarly configured interfaces are grouped based on hardware, administrative, and port parameter constraints. For example, LACP groups the interfaces with the same speed, duplex mode, native VLAN, VLAN range, and trunking status and type. After grouping the links into an EtherChannel, LACP adds the group to the spanning tree as a single switch port.

LACP Modes

Table 28-2 shows the user-configurable EtherChannel LACP modes for the **channel-group** interface configuration command.

Table 28-2 EtherChannel LACP Modes

Mode	Description
active	Places an interface into an active negotiating state in which the interface starts negotiations with other interfaces by sending LACP packets.
passive	Places an interface into a passive negotiating state in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. This setting minimizes the transmission of LACP packets.
on	Forces the interface to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode.

Both the **active** and **passive** LACP modes enable interfaces to negotiate with partner interfaces to determine if they can form an EtherChannel based on criteria such as interface speed and, for Layer 2 EtherChannels, trunking state and VLAN numbers.

Interfaces can form an EtherChannel when they are in different LACP modes as long as the modes are compatible. For example:

- An interface in the **active** mode can form an EtherChannel with another interface that is in the **active** or **passive** mode.
- An interface in the **passive** mode cannot form an EtherChannel with another interface that is also in the **passive** mode because neither interface starts LACP negotiation.

LACP Interaction with Other Features

The DTP and the CDP send and receive packets over the physical interfaces in the EtherChannel. Trunk ports send and receive LACP PDUs on the lowest numbered VLAN.

In Layer 2 EtherChannels, the first port in the channel that comes up provides its MAC address to the EtherChannel. If this port is removed from the bundle, one of the remaining ports in the bundle provides its MAC address to the EtherChannel.

LACP sends and receives LACP PDUs only from interfaces that are up and have LACP enabled for the active or passive mode.

Load Balancing and Forwarding Methods

EtherChannel balances the traffic load across the links in a channel by reducing part of the binary pattern formed from the addresses in the frame to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The selected mode applies to all EtherChannels configured on the switch. You configure the load balancing and forwarding method by using the **port-channel load-balance** global configuration command.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel.

With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the destination host's MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

With source-and-destination MAC address forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on both the source and destination MAC addresses. This forwarding method, a combination source-MAC and destination-MAC address forwarding methods of load distribution, can be used if it is not clear whether source-MAC or destination-MAC address forwarding is better suited on a particular switch. With source-and-destination MAC-address forwarding, packets sent from host A to host B, host A to host C, and host C to host B could all use different ports in the channel.

With source-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the source-IP address of the incoming packet. Therefore, to provide load-balancing, packets from different IP addresses use different ports in the channel, but packets from the same IP address use the same port in the channel.

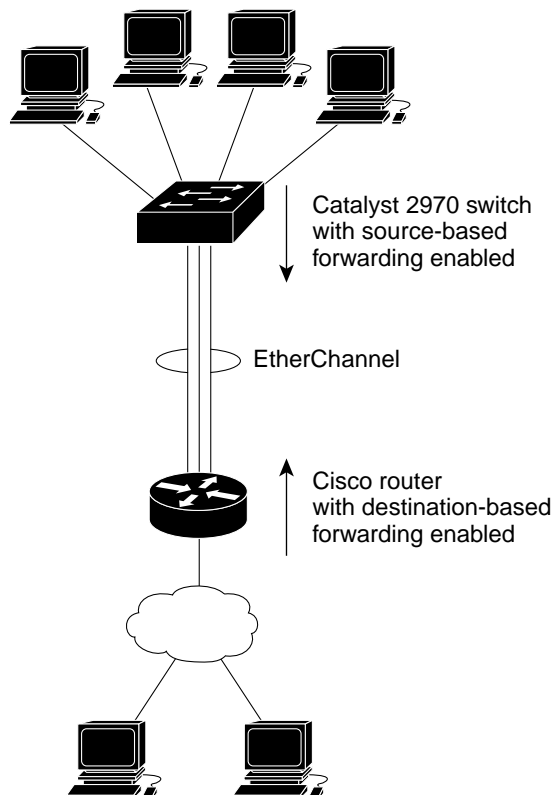
With destination-IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on the destination-IP address of the incoming packet. Therefore, to provide load-balancing, packets from the same IP source address sent to different IP destination addresses could be sent on different ports in the channel. But packets sent from different source IP addresses to the same destination IP address are always sent on the same port in the channel.

With source-and-destination IP address-based forwarding, when packets are forwarded to an EtherChannel, they are distributed across the ports in the EtherChannel based on both the source and destination IP addresses of the incoming packet. This forwarding method, a combination of source-IP and destination-IP address-based forwarding, can be used if it is not clear whether source-IP or destination-IP address-based forwarding is better suited on a particular switch. In this method, packets sent from the IP address A to IP address B, from IP address A to IP address C, and from IP address C to IP address B could all use different ports in the channel.

Different load-balancing methods have different advantages, and the choice of a particular load-balancing method should be based on the position of the switch in the network and the kind of traffic that needs to be load-distributed. In [Figure 28-3](#), an EtherChannel of four workstations communicates with a router. Because the router is a single-MAC-address device, source-based forwarding on the switch EtherChannel ensures that the switch uses all available bandwidth to the router. The router is configured for destination-based forwarding because the large number of workstations ensures that the traffic is evenly distributed from the router EtherChannel.

Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel is going only to a single MAC address, using the destination-MAC address always chooses the same link in the channel. Using source addresses or IP addresses might result in better load balancing.

Figure 28-3 Load Distribution and Forwarding Methods



90569

Configuring EtherChannels

These sections describe how to configure EtherChannel on Layer 2 interfaces:

- [Default EtherChannel Configuration, page 28-9](#)
- [EtherChannel Configuration Guidelines, page 28-9](#)
- [Configuring Layer 2 EtherChannels, page 28-10](#) (required)
- [Configuring EtherChannel Load Balancing, page 28-12](#) (optional)
- [Configuring the PAgP Learn Method and Priority, page 28-13](#) (optional)
- [Configuring LACP Hot-Standby Ports, page 28-15](#) (optional)



Note

Make sure that the interfaces are correctly configured. For more information, see the [“EtherChannel Configuration Guidelines”](#) section on page 28-9.



Note

After you configure an EtherChannel, configuration changes applied to the port-channel interface apply to all the physical interfaces assigned to the port-channel interface, and configuration changes applied to the physical interface affect only the interface where you apply the configuration.

Default EtherChannel Configuration

Table 28-3 shows the default EtherChannel configuration.

Table 28-3 *Default EtherChannel Configuration*

Feature	Default Setting
Channel groups	None assigned.
Port-channel logical interface	None defined.
PAgP mode	No default.
PAgP learn method	Aggregate-port learning on all interfaces.
PAgP priority	128 on all interfaces.
LACP mode	No default.
LACP learn method	Aggregate-port learning on all interfaces.
LACP port priority	32768 on all interfaces.
LACP system priority	32768.
LACP system ID	LACP system priority and the switch MAC address.
Load balancing	Load distribution on the switch is based on the source-MAC address of the incoming packet.

EtherChannel Configuration Guidelines

If improperly configured, some EtherChannel interfaces are automatically disabled to avoid network loops and other problems. Follow these guidelines to avoid configuration problems:

- More than 12 EtherChannels cannot be configured on a Catalyst 2970 switch.
- Configure a PAgP EtherChannel with up to eight Ethernet interfaces of the same type.
- Configure a LACP EtherChannel with up to 16 Ethernet interfaces of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
- Configure all interfaces in an EtherChannel to operate at the same speeds and duplex modes.
- Enable all interfaces in an EtherChannel. An interface in an EtherChannel that is disabled by using the **shutdown** interface configuration command is treated as a link failure, and its traffic is transferred to one of the remaining interfaces in the EtherChannel.
- When a group is first created, all ports follow the parameters set for the first port to be added to the group. If you change the configuration of one of these parameters, you must also make the changes to all ports in the group:
 - Allowed-VLAN list
 - Spanning-tree path cost for each VLAN
 - Spanning-tree port priority for each VLAN
 - Spanning-tree Port Fast setting
- Do not configure a port to be a member of more than one EtherChannel group.

- Do not configure an EtherChannel in both the PAgP and LACP modes. EtherChannel groups running PAgP and LACP can coexist on the same switch. Individual EtherChannel groups can run either PAgP or LACP, but they cannot interoperate.
- Do not configure a Switched Port Analyzer (SPAN) destination as part of an EtherChannel.
- Do not configure a secure port as part of an EtherChannel or the reverse.
- Do not configure a port that is an active member of an EtherChannel as an 802.1X port. If 802.1X is enabled on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
- For Layer 2 EtherChannels:
 - Assign all interfaces in the EtherChannel to the same VLAN, or configure them as trunks. Interfaces with different native VLANs cannot form an EtherChannel.
 - If you configure an EtherChannel from trunk interfaces, verify that the trunking mode (ISL or 802.1Q) is the same on all the trunks. Inconsistent trunk modes on EtherChannel interfaces can have unexpected results.
 - An EtherChannel supports the same allowed range of VLANs on all the interfaces in a trunking Layer 2 EtherChannel. If the allowed range of VLANs is not the same, the interfaces do not form an EtherChannel even when PAgP is set to the **auto** or **desirable** mode.
 - Interfaces with different spanning-tree path costs can form an EtherChannel if they are otherwise compatibly configured. Setting different spanning-tree path costs does not, by itself, make interfaces incompatible for the formation of an EtherChannel.

Configuring Layer 2 EtherChannels

You configure Layer 2 EtherChannels by assigning interfaces to a channel group with the **channel-group** interface configuration command. This command automatically creates the port-channel logical interface.

Beginning in privileged EXEC mode, follow these steps to assign a Layer 2 Ethernet interface to a Layer 2 EtherChannel. This procedure is required.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify a physical interface to configure. Valid interfaces include physical interfaces. For a PAgP EtherChannel, you can configure up to eight interfaces of the same type and speed for the same group. For a LACP EtherChannel, you can configure up to 16 Ethernet interfaces of the same type. Up to eight ports can be active, and up to eight ports can be in standby mode.
Step 3	switchport mode { access trunk } switchport access vlan <i>vlan-id</i>	Assign all interfaces as static-access ports in the same VLAN, or configure them as trunks. If you configure the interface as a static-access port, assign it to only one VLAN. The range is 1 to 4094.

	Command	Purpose
Step 4	channel-group <i>channel-group-number</i> mode { auto [non-silent] desirable [non-silent] on } { active passive }	<p>Assign the interface to a channel group, and specify the PAgP or the LACP mode.</p> <p>For <i>channel-group-number</i>, the range is 1 to 12.</p> <p>For mode, select one of these keywords:</p> <ul style="list-style-type: none"> • auto—Enables PAgP only if a PAgP device is detected. It places an interface into a passive negotiating state, in which the interface responds to PAgP packets it receives but does not start PAgP packet negotiation. • desirable—Unconditionally enables PAgP. It places an interface into an active negotiating state, in which the interface starts negotiations with other interfaces by sending PAgP packets. • on—Forces the interface to channel without PAgP or LACP. With the on mode, a usable EtherChannel exists only when an interface group in the on mode is connected to another interface group in the on mode. • non-silent—(Optional) If your switch is connected to a partner that is PAgP-capable, configure the switch interface for nonsilent operation when the interface is in the auto or desirable mode. If you do not specify non-silent, silent is assumed. The silent setting is for connections to file servers or packet analyzers. This setting allows PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission. • active—Enables LACP only if a LACP device is detected. It places an interface into an active negotiating state in which the interface starts negotiations with other interfaces by sending LACP packets. • passive—Enables LACP on an interface and places it into a passive negotiating state in which the interface responds to LACP packets that it receives, but does not start LACP packet negotiation. <p>For information on compatible modes for the switch and its partner, see the “PAgP Modes” section on page 28-4 and the “LACP Modes” section on page 28-6.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an interface from the EtherChannel group, use the **no channel-group** interface configuration command.

This example shows how to configure an EtherChannel. It assigns Gigabit Ethernet interfaces 0/3 and 0/4 as static-access ports in VLAN 10 to channel 5 with the PAgP mode **desirable**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/3 -4
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode desirable non-silent
Switch(config-if-range)# end
```

This example shows how to configure an EtherChannel. It assigns Gigabit Ethernet interfaces 0/3 and 0/4 as static-access ports in VLAN 10 to channel 5 with the LACP mode **active**:

```
Switch# configure terminal
Switch(config)# interface range gigabitethernet0/3 -4
Switch(config-if-range)# switchport mode access
Switch(config-if-range)# switchport access vlan 10
Switch(config-if-range)# channel-group 5 mode active
Switch(config-if-range)# end
```

Configuring EtherChannel Load Balancing

This section describes how to configure EtherChannel load balancing by using source-based or destination-based forwarding methods. For more information, see the [“Load Balancing and Forwarding Methods” section on page 28-6](#).

Beginning in privileged EXEC mode, follow these steps to configure EtherChannel load balancing. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	port-channel load-balance {dst-ip dst-mac src-dst-ip src-dst-mac src-ip src-mac}	<p>Configure an EtherChannel load-balancing method.</p> <p>The default is src-mac.</p> <p>Select one of these keywords to determine the load-distribution method:</p> <ul style="list-style-type: none"> • dst-ip—Load distribution is based on the destination-host IP address. • dst-mac—Load distribution is based on the destination-host MAC address of the incoming packet. • src-dst-ip—Load distribution is based on the source-and-destination host-IP address. • src-dst-mac—Load distribution is based on the source-and-destination host-MAC address. • src-ip—Load distribution is based on the source-host IP address. • src-mac—Load distribution is based on the source-MAC address of the incoming packet.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show etherchannel load-balance	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return EtherChannel load balancing to the default configuration, use the **no port-channel load-balance** global configuration command.

Configuring the PAgP Learn Method and Priority

Network devices are classified as PAgP physical learners or aggregate-port learners. A device is a physical learner if it learns addresses by physical ports and directs transmissions based on that knowledge. A device is an aggregate-port learner if it learns addresses by aggregate (logical) ports. The learn method must be configured the same at both ends of the link.

When a device and its partner are both aggregate-port learners, they learn the address on the logical port-channel. The device sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.

PAgP cannot automatically detect when the partner device is a physical learner and when the local device is an aggregate-port learner. Therefore, you must manually set the learning method on the local device to learn addresses by physical ports. You also must set the load-distribution method to source-based distribution, so that any given source MAC address is always sent on the same physical port.

You also can configure a single interface within the group for all transmissions and use other interfaces for hot standby. The unused interfaces in the group can be swapped into operation in just a few seconds if the selected single interface loses hardware-signal detection. You can configure which interface is always selected for packet transmission by changing its priority with the **pagp port-priority** interface configuration command. The higher the priority, the more likely that the port will be selected.



Note

The Catalyst 2970 switch supports address learning only on aggregate ports even though the **physical-port** keyword is provided in the CLI. The **pagp learn-method** command and the **pagp port-priority** command have no effect on the switch hardware, but they are required for PAgP interoperability with devices that only support address learning by physical ports, such as the Catalyst 1900 switch.

When the link partner to the Catalyst 2970 switch is a physical learner (such as a Catalyst 1900 series switch), we recommend that you configure the Catalyst 2970 switch as a physical-port learner by using the **pagp learn-method physical-port** interface configuration command. Set the load-distribution method based on the source MAC address by using the **port-channel load-balance src-mac** global configuration command. The switch then sends packets to the Catalyst 1900 switch using the same interface in the EtherChannel from which it learned the source address. Use the **pagp learn-method** command only in this situation.

Beginning in privileged EXEC mode, follow these steps to configure your switch as a PAgP physical-port learner and to adjust the priority so that the same port in the bundle is selected for sending packets. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface for transmission.
Step 3	pagp learn-method physical-port	<p>Select the PAgP learning method.</p> <p>By default, aggregation-port learning is selected, which means the switch sends packets to the source by using any of the interfaces in the EtherChannel. With aggregate-port learning, it is not important on which physical port the packet arrives.</p> <p>Select physical-port to connect with another switch that is a physical learner. Make sure to configure the port-channel load-balance global configuration command to src-mac as described in the “Configuring EtherChannel Load Balancing” section on page 28-12.</p> <p>The learning method must be configured the same at both ends of the link.</p>
Step 4	pagp port-priority <i>priority</i>	<p>Assign a priority so that the selected interface is chosen for packet transmission.</p> <p>For <i>priority</i>, the range is 0 to 255. The default is 128. The higher the priority, the more likely that the interface will be used for PAgP transmission.</p>
Step 5	end	Return to privileged EXEC mode.
Step 6	show running-config or show pagp <i>channel-group-number</i> internal	Verify your entries.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the priority to its default setting, use the **no pagp port-priority** interface configuration command. To return the learning method to its default setting, use the **no pagp learn-method** interface configuration command.

Configuring LACP Hot-Standby Ports

When enabled, LACP tries to configure the maximum number of LACP-compatible ports in a channel, up to a maximum of 16 ports. Only eight LACP links can be active at one time. The software places any additional links in a hot-standby mode. If one of the active links becomes inactive, a link that is in the hot-standby mode becomes active in its place.

If you configure more than eight links for an EtherChannel group, the software automatically determines which of the hot-standby ports to make active based on the LACP priority. The software assigns to every link between systems that operate LACP a unique priority made up of these elements (in priority order):

- LACP system priority
- System ID (a combination of the LACP system priority and the switch MAC address)
- LACP port priority
- Port number

In priority comparisons, numerically lower values have higher priority. The priority determines which ports should be put in standby mode when there is a hardware limitation that prevents all compatible ports from aggregating.

Ports are considered for active use in aggregation in link-priority order starting with the port attached to the highest priority link. Each port is selected for active use if the preceding higher priority selections can also be maintained. Otherwise, the port is selected for standby mode.

You can change the default values of the LACP system priority and the LACP port priority to affect how the software selects active and standby links. For more information, see the [“Configuring the LACP System Priority” section on page 28-15](#) and the [“Configuring the LACP Port Priority” section on page 28-16](#).

Configuring the LACP System Priority

You can configure the system priority for all of the EtherChannels that are enabled for LACP by using the **lacp system-priority** global configuration command. You cannot configure a system priority for each LACP-configured channel. By changing this value from the default, you can affect how the software selects active and standby links.

You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an *H* port-state flag).

Beginning in privileged EXEC mode, follow these steps to configure the LACP system priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	lacp system-priority <i>priority</i>	Configure the LACP system priority. For <i>priority</i> , the range is 1 to 65535. The default is 32768. The lower the value, the higher the system priority.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show running-config or show lacp sys-id	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP system priority to the default value, use the **no lacp system-priority** global configuration command.

Configuring the LACP Port Priority

By default, all ports use the same port priority. If the local system has a lower value for the system priority and the system ID than the remote system, you can affect which of the hot-standby links become active first by changing the port priority of LACP EtherChannel ports to a lower value than the default. The hot-standby ports that have lower port numbers become active in the channel first. You can use the **show etherchannel summary** privileged EXEC command to see which ports are in the hot-standby mode (denoted with an *H* port-state flag).



Note

If LACP is not able to aggregate all the ports that are compatible (for example, the remote system might have more restrictive hardware limitations), all the ports that cannot be actively included in the EtherChannel are put in the hot-standby state and are used only if one of the channeled ports fails.

Beginning in privileged EXEC mode, follow these steps to configure the LACP port priority. This procedure is optional.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to be configured.
Step 3	lacp port-priority <i>priority</i>	Configure the LACP port priority. For <i>priority</i> , the range is 1 to 65535. The is 32768. The lower the value, the more likely that the interface will be used for LACP transmission.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config or show lacp [<i>channel-group-number</i>] internal	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return the LACP port priority to the default value, use the **no lacp port-priority** interface configuration command.

Displaying EtherChannel, PAgP, and LACP Status

To display EtherChannel, PAgP, and LACP status information, use the privileged EXEC commands described in [Table 28-4](#):

Table 28-4 *Commands for Displaying EtherChannel, PAgP, and LACP Status*

Command	Description
show etherchannel [<i>channel-group-number</i> { detail port port-channel protocol summary }] { detail load-balance port port-channel protocol summary }	Displays EtherChannel information in a brief, detailed, and one-line summary form. Also displays the load-balance or frame-distribution scheme, port, port-channel, and protocol information.
show pagp [<i>channel-group-number</i>] { counters internal neighbor }	Displays PAgP information such as traffic information, the internal PAgP configuration, and neighbor information.
show lacp [<i>channel-group-number</i>] { counters internal neighbor }	Displays LACP information such as traffic information, the internal LACP configuration, and neighbor information.

You can clear PAgP channel-group information and traffic counters by using the **clear pagp** {*channel-group-number* **counters** | **counters**} privileged EXEC command.

You can clear LACP channel-group information and traffic counters by using the **clear lacp** {*channel-group-number* **counters** | **counters**} privileged EXEC command.

For detailed information about the fields in the displays, refer to the command reference for this release.



Troubleshooting

This chapter describes how to identify and resolve software problems related to the Cisco IOS software on the Catalyst 2970 switch. Depending on the nature of the problem, you can use the command-line interface (CLI) or the Cluster Management Suite (CMS) to identify and solve problems.

Additional troubleshooting information, such as LED descriptions, is provided in the hardware installation guide.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the command reference for this release and the *Cisco IOS Command Summary for Release 12.1*.

This chapter consists of these sections:

- [Recovering from Corrupted Software By Using the XMODEM Protocol, page 29-2](#)
- [Recovering from a Lost or Forgotten Password, page 29-4](#)
- [Recovering from a Command Switch Failure, page 29-8](#)
- [Recovering from Lost Cluster Member Connectivity, page 29-11](#)



Note

Recovery procedures require that you have physical access to the switch.

- [Preventing Autonegotiation Mismatches, page 29-12](#)
- [SFP Module Security and Identification, page 29-12](#)
- [Using Ping, page 29-13](#)
- [Using Layer 2 Traceroute, page 29-14](#)
- [Using IP Traceroute, page 29-16](#)
- [Using Debug Commands, page 29-17](#)
- [Using the show platform forward Command, page 29-19](#)
- [Using the crashinfo File, page 29-21](#)

Recovering from Corrupted Software By Using the XMODEM Protocol

Switch software can be corrupted during an upgrade, by downloading the wrong file to the switch, and by deleting the image file. In all of these cases, the switch does not pass the power-on self-test (POST), and there is no connectivity.

This procedure uses the XMODEM Protocol to recover from a corrupt or wrong image file. There are many software packages that support the XMODEM Protocol, and this procedure is largely dependent on the emulation software you are using.

This recovery procedure requires that you have physical access to the switch.

Step 1 From your PC, download the software image tar file (*image_filename.tar*) from Cisco.com.

The Cisco IOS image is stored as a bin file in a directory in the tar file. For information about locating the software image files on Cisco.com, refer to the release notes.

Step 2 Extract the bin file from the tar file.

- If you are using Windows, use a zip program that is capable of reading a tar file. Use the zip program to navigate to and extract the bin file.
- If you are using UNIX, follow these steps:
 1. Display the contents of the tar file by using the **tar -tvf <image_filename.tar>** UNIX command.

```
switch% tar -tvf image_filename.tar
drwxr-xr-x 9658/25      0 Apr 21 13:24 2003 c2970-i612-mz.121.11-AX/
drwxr-xr-x 9658/25      0 Apr 18 16:39 2003 c2970-i612-mz.121.11-AX/html/
-rw-r--r-- 9658/25    4005 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/homepage.htm
-rw-r--r-- 9658/25    1392 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/not_supported.html
-rw-r--r-- 9658/25    9448 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/common.js
-rw-r--r-- 9658/25    22152 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/cms_splash.gif
-rw-r--r-- 9658/25    1211 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/cms_13.html
-rw-r--r-- 9658/25    2823 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/cluster.html
-rw-r--r-- 9658/25    4195 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/Redirect.jar
-rw-r--r-- 9658/25   14984 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/mono_disc.sgz
-rw-r--r-- 9658/25  1329516 Apr 18 15:56 2003 c2970-i612-mz.121.11-AX/html/CMS.sgz
-rw-r--r-- 9658/25  140105 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/images.sgz
-rw-r--r-- 9658/25  213848 Apr 18 15:56 2003 c2970-i612-mz.121.11-AX/html/help.sgz
-rw-r--r-- 9658/25  135599 Apr 18 15:56 2003
c2970-i612-mz.121.11-AX/html/CiscoChartPanel.sgz
-rwxr-xr-x 9658/25    58862 Apr 18 18:45 2003
c2970-i612-mz.121.11-AX/html/cms_boot.jar
-rw-r--r-- 9658/25  2928176 Apr 21 12:01 2003
c2970-i612-mz.121.11-AX/c2970-i612-mz.121.11-AX.bin
-rw-r--r-- 9658/25     320 Apr 21 13:24 2003 c2970-i612-mz.121.11-AX/info
-rw-r--r-- 9658/25     100 Apr 18 16:59 2003 info
```

2. Locate the bin file and extract it by using the **tar -xvf <image_filename.tar> <image_filename.bin>** UNIX command.

```
switch% tar -xvf image_filename.tar image_filename.bin
x c2970-i612-mz.121.11-AX/c2970-i612-mz.121.11-AX.bin, 2928176 bytes, 5720 tape
blocks
```

3. Verify that the bin file was extracted by using the **ls -l <image_filename.bin>** UNIX command.

```
switch% ls -l image_filename.bin
-rw-r--r-- 1 boba          2928176 Apr 21 12:01
c2970-i612-mz.121.11-AX/c2970-i612-mz.121.11-AX.bin
```

Step 3 Connect your PC with terminal-emulation software supporting the XMODEM Protocol to the switch console port.

Step 4 Set the line speed on the emulation software to 9600 baud.

Step 5 Unplug the switch power cord.

Step 6 Press the **Mode** button, and at the same time, reconnect the power cord to the switch.

You can release the **Mode** button a second or two after the LED above port 1 goes off. Several lines of information about the software appear along with instructions:

```
The system has been interrupted prior to initializing the flash file system. The following
commands will initialize the flash file system, and finish loading the operating system
software#
```

```
flash_init
load_helper
boot
```

Step 7 Initialize the Flash file system:

```
switch: flash_init
```

Step 8 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

Step 9 Load any helper files:

```
switch: load_helper
```

Step 10 Start the file transfer by using the XMODEM protocol.

```
switch: copy xmodem: flash:image_filename.bin
```

Step 11 After the XMODEM request appears, use the appropriate command on the terminal-emulation software to start the transfer and to copy the software image into Flash memory.

Step 12 Boot the newly downloaded Cisco IOS image.

```
switch:boot flash:image_filename.bin
```

Step 13 Use the **archive download-sw** privileged EXEC command to download the software image to the switch.

Step 14 Use the **reload** privileged EXEC command to restart the switch and to verify that the new software image is operating properly.

Step 15 Delete the `flash:image_filename.bin` file from the switch.

Recovering from a Lost or Forgotten Password

The default configuration for the switch allows an end user with physical access to the switch to recover from a lost password by interrupting the boot process during power-on and by entering a new password. These recovery procedures require that you have physical access to the switch.

**Note**

On these switches, a system administrator can disable some of the functionality of this feature by allowing an end user to reset a password only by agreeing to return to the default configuration. If you are an end user trying to reset a password when password recovery has been disabled, a status message shows this during the recovery process.

This section describes how to recover a forgotten or lost switch password. It provides two solutions:

- [Procedure with Password Recovery Enabled, page 29-5](#)
- [Procedure with Password Recovery Disabled, page 29-6](#)

You enable or disable password recovery by using the **service password-recovery** global configuration command. Follow the steps in this procedure if you have forgotten or lost the switch password.

Step 1 Connect a terminal or PC with terminal-emulation software to the switch console port.

Step 2 Set the line speed on the emulation software to 9600 baud.

Step 3 Power off the switch.

Step 4 Press the **Mode** button, and at the same time, reconnect the power cord to the switch.

You can release the **Mode** button a second or two after the LED above port 1 turns off. Several lines of information about the software appear with instructions, informing you if the password recovery procedure has been disabled or not.

- If you see a message that begins with this:

```
The system has been interrupted prior to initializing the flash file system. The
following commands will initialize the flash file system
```

proceed to the [“Procedure with Password Recovery Enabled” section on page 29-5](#), and follow the steps.

- If you see a message that begins with this:

```
The password-recovery mechanism has been triggered, but is currently disabled.
```

proceed to the [“Procedure with Password Recovery Disabled” section on page 29-6](#), and follow the steps.

Step 5 After recovering the password, reload the switch:

```
Switch> reload
Proceed with reload? [confirm] y
```

Procedure with Password Recovery Enabled

If the password-recovery mechanism is enabled, this message appears:

The system has been interrupted prior to initializing the flash file system. The following commands will initialize the flash file system, and finish loading the operating system software:

```
flash_init
load_helper
boot
```

Step 1 Initialize the Flash file system:

```
switch: flash_init
```

Step 2 If you had set the console port speed to anything other than 9600, it has been reset to that particular speed. Change the emulation software line speed to match that of the switch console port.

Step 3 Load any helper files:

```
switch: load_helper
```

Step 4 Display the contents of Flash memory:

```
switch: dir flash:
```

The switch file system appears:

Directory of flash:

```
 13 drwx      192  Mar 01 1993 22:30:48
 11 -rwx      5825 Mar 01 1993 22:31:59 config.text
 18 -rwx      720  Mar 01 1993 02:21:30 vlan.dat
```

```
16128000 bytes total (10003456 bytes free)
```

Step 5 Rename the configuration file to config.text.old.

This file contains the password definition.

```
switch: rename flash:config.text flash:config.text.old
```

Step 6 Boot the system:

```
switch: boot
```

You are prompted to start the setup program. Enter **N** at the prompt:

```
Continue with the configuration dialog? [yes/no]: N
```

Step 7 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
```

Step 8 Rename the configuration file to its original name:

```
Switch# rename flash:config.text.old flash:config.text
```

Step 9 Copy the configuration file into memory:

```
Switch# copy flash:config.text system:running-config
Source filename [config.text]?
Destination filename [running-config]?
```

Press **Return** in response to the confirmation prompts.

The configuration file is now reloaded, and you can change the password.

Step 10 Enter global configuration mode:

```
Switch# configure terminal
```

Step 11 Change the password:

```
Switch (config)# enable secret password
```

The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

Step 12 Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

Step 13 Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.

**Note**

This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan *vlan-id*** global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

Step 14 Reload the switch:

```
Switch# reload
```

Procedure with Password Recovery Disabled

If the password-recovery mechanism is disabled, this message appears:

```
The password-recovery mechanism has been triggered, but
is currently disabled. Access to the boot loader prompt
through the password-recovery mechanism is disallowed at
this point. However, if you agree to let the system be
reset back to the default system configuration, access
to the boot loader prompt can still be allowed.
```

```
Would you like to reset the system back to the default configuration (y/n)?
```

**Caution**

Returning the switch to the default configuration results in the loss of all existing configurations. We recommend that you contact your system administrator to verify if there are backup switch and VLAN configuration files.

- If you enter **n** (no), the normal boot process continues as if the **Mode** button had not been pressed; you cannot access the boot loader prompt, and you cannot enter a new password. You see the message:

```
Press Enter to continue.....
```

- If you enter **y** (yes), the configuration file in Flash memory and the VLAN database file are deleted. When the default configuration loads, you can reset the password.

Step 1 Elect to continue with password recovery and lose the existing configuration:

```
Would you like to reset the system back to the default configuration (y/n)? Y
```

Step 2 Load any helper files:

```
Switch: load_helper
```

Step 3 Display the contents of Flash memory:

```
switch: dir flash:
```

The switch file system appears:

```
Directory of flash:
13  drwx          192   Mar 01 1993 22:30:48

16128000 bytes total (10003456 bytes free)
```

Step 4 Boot the system:

```
Switch: boot
```

You are prompted to start the setup program. To continue with password recovery, enter **N** at the prompt:

```
Continue with the configuration dialog? [yes/no]: N
```

Step 5 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
```

Step 6 Enter global configuration mode:

```
Switch# configure terminal
```

Step 7 Change the password:

```
Switch (config)# enable secret password
```

The secret password can be from 1 to 25 alphanumeric characters, can start with a number, is case sensitive, and allows spaces but ignores leading spaces.

Step 8 Return to privileged EXEC mode:

```
Switch (config)# exit
Switch#
```

- Step 9 Write the running configuration to the startup configuration file:

```
Switch# copy running-config startup-config
```

The new password is now in the startup configuration.

**Note**

This procedure is likely to leave your switch virtual interface in a shutdown state. You can see which interface is in this state by entering the **show running-config** privileged EXEC command. To re-enable the interface, enter the **interface vlan vlan-id** global configuration command, and specify the VLAN ID of the shutdown interface. With the switch in interface configuration mode, enter the **no shutdown** command.

- Step 10 You must now reconfigure the switch. If the system administrator has the backup switch and VLAN configuration files available, you should use those.

Recovering from a Command Switch Failure

This section describes how to recover from a failed command switch. You can configure a redundant command switch group by using the Hot Standby Router Protocol (HSRP). For more information, see [Chapter 5, “Clustering Switches.”](#)

**Note**

HSRP is the preferred method for supplying redundancy to a cluster.

If you have not configured a standby command switch, and your command switch loses power or fails in some other way, management contact with the member switches is lost, and you must install a new command switch. However, connectivity between switches that are still connected is not affected, and the member switches forward packets as usual. You can manage the members as standalone switches through the console port or, if they have IP addresses, through the other management interfaces.

You can prepare for a command switch failure by assigning an IP address to a member switch or another switch that is command-capable, making a note of the command-switch password, and cabling your cluster to provide redundant connectivity between the member switches and the replacement command switch. This section describes two solutions for replacing a failed command switch:

- [Replacing a Failed Command Switch with a Cluster Member, page 29-8](#)
- [Replacing a Failed Command Switch with Another Switch, page 29-10](#)

These recovery procedures require that you have physical access to the switch.

For information on command-capable switches, refer to the release notes.

Replacing a Failed Command Switch with a Cluster Member

To replace a failed command switch with a command-capable member in the same cluster, follow these steps:

- Step 1 Disconnect the command switch from the member switches, and physically remove it from the cluster.
- Step 2 Insert the member switch in place of the failed command switch, and duplicate its connections to the cluster members.

- Step 3** Start a CLI session on the new command switch.
- You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, refer to the switch hardware installation guide.
- Step 4** At the switch prompt, enter privileged EXEC mode:
- ```
Switch> enable
Switch#
```
- Step 5** Enter the password of the *failed command switch*.
- Step 6** Enter global configuration mode.
- ```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
```
- Step 7** Remove the member switch from the cluster.
- ```
Switch(config)# no cluster commander-address
```
- Step 8** Return to privileged EXEC mode.
- ```
Switch(config)# end
Switch#
```
- Step 9** Use the setup program to configure the switch IP information. This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.
- ```
Switch# setup
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:
```
- Step 10** Enter **Y** at the first prompt.
- The prompts in the setup program vary depending on the member switch you selected to be the command switch:
- ```
Continue with configuration dialog? [yes/no]: y
or
Configuring global parameters:
```
- If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.
- Step 11** Respond to the questions in the setup program.
- When prompted for the host name, recall that on a command switch, the host name is limited to 28 characters; on a member switch to 31 characters. Do not use *-n*, where *n* is a number, as the last characters in a host name for any switch.
- When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

- Step 12 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.
 - Step 13 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.
 - Step 14 When prompted, assign a name to the cluster, and press **Return**.
The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.
 - Step 15 After the initial configuration displays, verify that the addresses are correct.
 - Step 16 If the displayed information is correct, enter **Y**, and press **Return**.
If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.
 - Step 17 Start your browser, and enter the IP address of the new command switch.
 - Step 18 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.
-

Replacing a Failed Command Switch with Another Switch

To replace a failed command switch with a switch that is command-capable but not part of the cluster, follow these steps:

- Step 1 Insert the new switch in place of the failed command switch, and duplicate its connections to the cluster members.
- Step 2 Start a CLI session on the new command switch.
You can access the CLI by using the console port or, if an IP address has been assigned to the switch, by using Telnet. For details about using the console port, refer to the switch hardware installation guide.
- Step 3 At the switch prompt, enter privileged EXEC mode:

```
Switch> enable
Switch#
```
- Step 4 Enter the password of the *failed command switch*.
- Step 5 Use the setup program to configure the switch IP information.

This program prompts you for IP address information and passwords. From privileged EXEC mode, enter **setup**, and press **Return**.

```
Switch# setup
--- System Configuration Dialog ---
Continue with configuration dialog? [yes/no]: y

At any point you may enter a question mark '?' for help.
Use ctrl-c to abort configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Basic management setup configures only enough connectivity
for management of the system, extended setup will ask you
to configure each interface on the system

Would you like to enter basic management setup? [yes/no]:
```

Step 6 Enter **Y** at the first prompt.

The prompts in the setup program vary depending on the switch you selected to be the command switch:

```
Continue with configuration dialog? [yes/no]: y
```

or

```
Configuring global parameters:
```

If this prompt does not appear, enter **enable**, and press **Return**. Enter **setup**, and press **Return** to start the setup program.

Step 7 Respond to the questions in the setup program.

When prompted for the host name, recall that on a command switch, the host name is limited to 28 characters. Do not use *-n*, where *n* is a number, as the last characters in a host name for any switch.

When prompted for the Telnet (virtual terminal) password, recall that it can be from 1 to 25 alphanumeric characters, is case sensitive, allows spaces, but ignores leading spaces.

Step 8 When prompted for the **enable secret** and **enable** passwords, enter the passwords of the *failed command switch* again.

Step 9 When prompted, make sure to enable the switch as the cluster command switch, and press **Return**.

Step 10 When prompted, assign a name to the cluster, and press **Return**.

The cluster name can be 1 to 31 alphanumeric characters, dashes, or underscores.

Step 11 When the initial configuration displays, verify that the addresses are correct.

Step 12 If the displayed information is correct, enter **Y**, and press **Return**.

If this information is not correct, enter **N**, press **Return**, and begin again at Step 9.

Step 13 Start your browser, and enter the IP address of the new command switch.

Step 14 From the Cluster menu, select **Add to Cluster** to display a list of candidate switches to add to the cluster.

Recovering from Lost Cluster Member Connectivity

Some configurations can prevent the command switch from maintaining contact with member switches. If you are unable to maintain management contact with a member, and the member switch is forwarding packets normally, check for these conflicts:

- A member switch (Catalyst 3750, Catalyst 3550, Catalyst 3500 XL, Catalyst 2970, Catalyst 2950, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) cannot connect to the command switch through a port that is defined as a network port.
- Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 member switches must connect to the command switch through a port that belongs to the same management VLAN.
- A member switch (Catalyst 3750, Catalyst 3550, Catalyst 2970, Catalyst 2950, Catalyst 3500 XL, Catalyst 2900 XL, Catalyst 2820, and Catalyst 1900 switch) connected to the command switch through a secured port can lose connectivity if the port is disabled because of a security violation.

Preventing Autonegotiation Mismatches

The IEEE 802.3AB autonegotiation protocol manages the switch settings for speed (10 Mbps, 100 Mbps, and 1000 Mbps, excluding SFP module ports) and duplex (half or full). There are situations when this protocol can incorrectly align these settings, reducing performance. A mismatch occurs under these circumstances:

- A manually set speed or duplex parameter is different from the manually set speed or duplex parameter on the connected port.
- A port is set to autonegotiate, and the connected port is set to full duplex with no autonegotiation.

To maximize switch performance and ensure a link, follow one of these guidelines when changing the settings for duplex and speed:

- Let both ports autonegotiate both speed and duplex.
- Manually set the speed and duplex parameters for the ports on both ends of the connection.

**Note**

If a remote device does not autonegotiate, configure the duplex settings on the two ports to match. The speed parameter can adjust itself even if the connected port does not autonegotiate.

SFP Module Security and Identification

Cisco-approved small form-factor pluggable (SFP) modules have a serial EEPROM that contains the module serial number, the vendor name and ID, a unique security code, and cyclic redundancy check (CRC). When an SFP module is inserted in the switch, the switch software reads the EEPROM to verify the serial number, vendor name and vendor ID, and recompute the security code and CRC. If the serial number, the vendor name or vendor ID, the security code, or CRC is invalid, the software generates a security error message and places the interface in an error-disabled state.

**Note**

The security error message references the GBIC_SECURITY facility. The Catalyst 2970 switch supports SFP modules and does not support GBIC modules. Although the error message text refers to GBIC interfaces and modules, the security messages actually refer to the SFP modules and module interfaces. For more information about error messages, refer to the system message guide for this release.

If you are using a non-Cisco approved SFP module, remove the SFP module from the switch, and replace it with a Cisco-approved module. After inserting a Cisco-approved SFP module, use the **errdisable recovery cause gbic-invalid** global configuration command to verify the port status, and enter a time interval for recovering from the error-disabled state. After the elapsed interval, the switch brings the interface out of the error-disabled state and retries the operation. For more information about the **errdisable recovery** command, refer to the command reference for this release.

If the module is identified as a Cisco SFP module, but the system is unable to read vendor-data information to verify its accuracy, an SFP module error message is generated. In this case, you should remove and re-insert the SFP module. If it continues to fail, the SFP module might be defective.

Using Ping

This section consists of this information:

- [Understanding Ping, page 29-13](#)
- [Executing Ping, page 29-13](#)

Understanding Ping

The switch supports IP ping, which you can use to test connectivity to remote hosts. Ping sends an echo request packet to an address and waits for a reply. Ping returns one of these responses:

- Normal response—The normal response (*hostname* is alive) occurs in 1 to 10 seconds, depending on network traffic.
- Destination does not respond—If the host does not respond, a *no-answer* message is returned.
- Unknown host—If the host does not exist, an *unknown host* message is returned.
- Destination unreachable—If the default gateway cannot reach the specified network, a *destination-unreachable* message is returned.
- Network or host unreachable—If there is no entry in the route table for the host or network, a *network or host unreachable* message is returned.

Executing Ping

Beginning in privileged EXEC mode, use this command to ping another device on the network from the switch:

Command	Purpose
ping ip <i>host</i> <i>address</i>	Ping a remote host through IP or by supplying the host name or network address.



Note

Though other protocol keywords are available with the **ping** command, they are not supported in this release.

This example shows how to ping an IP host:

```
Switch# ping 172.20.52.3
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echoes to 172.20.52.3, timeout is 2 seconds:
```

```
!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

```
Switch#
```

[Table 29-1](#) describes the possible ping character output.

Table 29-1 Ping Output Display Characters

Character	Description
!	Each exclamation point means receipt of a reply.
.	Each period means the network server timed out while waiting for a reply.
U	A destination unreachable error PDU was received.
C	A congestion experienced packet was received.
I	User interrupted test.
?	Unknown packet type.
&	Packet lifetime exceeded.

To terminate a ping session, enter the escape sequence (**Ctrl-^ X** by default). You enter the default by simultaneously pressing and releasing the **Ctrl**, **Shift**, and **6** keys, and then pressing the **X** key.

Using Layer 2 Traceroute

This section describes this information:

- [Understanding Layer 2 Traceroute, page 29-14](#)
- [Usage Guidelines, page 29-14](#)
- [Displaying the Physical Path, page 29-15](#)

Understanding Layer 2 Traceroute

The Layer 2 traceroute feature allows the switch to identify the physical path that a packet takes from a source device to a destination device. Layer 2 traceroute supports only unicast source and destination MAC addresses. It determines the path by using the MAC address tables of the switches in the path. When the switch detects a device in the path that does not support Layer 2 traceroute, the switch continues to send Layer 2 trace queries and lets them time out.

The switch can only identify the path from the source device to the destination device. It cannot identify the path that a packet takes from source host to the source device or from the destination device to the destination host.

Usage Guidelines

These are the Layer 2 traceroute usage guidelines:

- Cisco Discovery Protocol (CDP) must be enabled on all the devices in the network. For Layer 2 traceroute to function properly, do not disable CDP.

For a list of switches that support Layer 2 traceroute, see the [“Usage Guidelines” section on page 29-14](#). If any devices in the physical path are transparent to CDP, the switch cannot identify the path through these devices.



Note For more information about enabling CDP, see [Chapter 20, “Configuring CDP.”](#)

- A switch is reachable from another switch when you can test connectivity by using the **ping** privileged EXEC command. All switches in the physical path must be reachable from each other.
- The maximum number of hops identified in the path is ten.
- You can enter the **traceroute mac** or the **traceroute mac ip** privileged EXEC command on a switch that is not in the physical path from the source device to the destination device. All switches in the path must be reachable from this switch.
- The **traceroute mac** command output shows the Layer 2 path only when the specified source and destination MAC addresses belong to the same VLAN. If you specify source and destination MAC addresses that belong to different VLANs, the Layer 2 path is not identified, and an error message appears.
- If you specify a multicast source or destination MAC address, the path is not identified, and an error message appears.
- If the source or destination MAC address belongs to multiple VLANs, you must specify the VLAN to which both the source and destination MAC addresses belong. If the VLAN is not specified, the path is not identified, and an error message appears.
- The **traceroute mac ip** command output shows the Layer 2 path when the specified source and destination IP addresses belong to the same subnet. When you specify the IP addresses, the switch uses the Address Resolution Protocol (ARP) to associate the IP addresses with the corresponding MAC addresses and the VLAN IDs.
 - If an ARP entry exists for the specified IP address, the switch uses the associated MAC address and identifies the physical path.
 - If an ARP entry does not exist, the switch sends an ARP query and tries to resolve the IP address. If the IP address is not resolved, the path is not identified, and an error message appears.
- When multiple devices are attached to one port through hubs (for example, multiple CDP neighbors are detected on a port), the Layer 2 traceroute feature is not supported. When more than one CDP neighbor is detected on a port, the Layer 2 path is not identified, and an error message appears.
- This feature is not supported in Token Ring VLANs.

Displaying the Physical Path

You can display physical path that a packet takes from a source device to a destination device by using one of these privileged EXEC commands:

- **tracetroute mac** [**interface** *interface-id*] {*source-mac-address*} [**interface** *interface-id*] {*destination-mac-address*} [**vlan** *vlan-id*] [**detail**]
- **tracetroute mac ip** {*source-ip-address* / *source-hostname*} {*destination-ip-address* / *destination-hostname*} [**detail**]

For more information, refer to the command reference for this release.

Using IP Traceroute

This section consists of this information:

- [Understanding IP Traceroute, page 29-16](#)
- [Executing IP Traceroute, page 29-16](#)

Understanding IP Traceroute

You can use IP traceroute to identify the path that packets take through the network on a hop-by-hop basis. The command output displays all network layer (Layer 3) devices, such as routers, that the traffic passes through on the way to the destination.

Your switches can participate as the source or destination of the **traceroute** privileged EXEC command and might or might not appear as a hop in the **traceroute** command output. If the switch is the destination of the traceroute, it is displayed as the final destination in the traceroute output. Intermediate switches do not show up in the traceroute output if they are only bridging the packet from one port to another within the same VLAN. However, if the intermediate switch is a multilayer switch that is routing a particular packet, this switch shows up as a hop in the traceroute output.

The **traceroute** privileged EXEC command uses the Time To Live (TTL) field in the IP header to cause routers and servers to generate specific return messages. Traceroute starts by sending a User Datagram Protocol (UDP) datagram to the destination host with the TTL field set to 1. If a router finds a TTL value of 1 or 0, it drops the datagram and sends back an Internet Control Message Protocol (ICMP) time-to-live-exceeded message to the sender. Traceroute determines the address of the first hop by examining the source address field of the ICMP time-to-live-exceeded message.

To identify the next hop, traceroute sends a UDP packet with a TTL value of 2. The first router decrements the TTL field by 1 and sends the datagram to the next router. The second router sees a TTL value of 1, discards the datagram, and returns the time-to-live-exceeded message to the source. This process continues until the TTL is incremented to a value large enough for the datagram to reach the destination host (or until the maximum TTL is reached).

To determine when a datagram reaches its destination, traceroute sets the UDP destination port number in the datagram to a very large value that the destination host is unlikely to be using. When a host receives a datagram destined to itself containing a destination port number that is unused locally, it sends an ICMP port unreachable error to the source. Because all errors except port unreachable errors come from intermediate hops, the receipt of a port unreachable error means this message was sent by the destination.

Executing IP Traceroute

Beginning in privileged EXEC mode, follow this step to trace the path packets take through the network:

Command	Purpose
traceroute ip <i>host</i>	Trace the path packets take through the network by using IP.

**Note**

Though other protocol keywords are available with the **traceroute** privileged EXEC command, they are not supported in this release.

This example shows how to perform a **traceroute** to an IP host:

```
Switch# traceroute ip 171.9.15.10

Type escape sequence to abort.
Tracing the route to 171.69.115.10

 0 172.2.52.1 0 msec 0 msec 4 msec
 1 172.2.1.203 12 msec 8 msec 0 msec
 2 171.9.16.6 4 msec 0 msec 0 msec
 3 171.9.4.5 0 msec 4 msec 0 msec
 4 171.9.121.34 0 msec 4 msec 4 msec
 5 171.9.15.9 120 msec 132 msec 128 msec
 6 171.9.15.10 132 msec 128 msec 128 msec
Switch#
```

The display shows the hop count, IP address of the router, and the round-trip time in milliseconds for each of the three probes that are sent.

Table 29-2 Traceroute Output Display Characters

Character	Description
*	The probe timed out.
?	Unknown packet type.
A	Administratively unreachable. Usually, this output means that an access list is blocking traffic.
H	Host unreachable.
N	Network unreachable.
P	Protocol unreachable.
Q	Source quench.
U	Port unreachable.

To terminate a trace in progress, enter the escape sequence (**Ctrl-^ X** by default). You enter the default by simultaneously pressing and releasing the **Ctrl**, **Shift**, and **6** keys, and then by pressing the **X** key.

Using Debug Commands

This section explains how you use **debug** commands to diagnose and resolve internetworking problems. It contains this information:

- [Enabling Debugging on a Specific Feature, page 29-18](#)
- [Enabling All-System Diagnostics, page 29-18](#)
- [Redirecting Debug and Error Message Output, page 29-19](#)



Caution

Because debugging output is assigned high priority in the CPU process, it can render the system unusable. For this reason, use **debug** commands only to troubleshoot specific problems or during troubleshooting sessions with Cisco technical support staff. It is best to use **debug** commands during periods of lower network traffic and fewer users. Debugging during these periods decreases the likelihood that increased **debug** command processing overhead will affect system use.

**Note**

For complete syntax and usage information for specific **debug** commands, refer to the command reference for this release.

Enabling Debugging on a Specific Feature

All **debug** commands are entered in privileged EXEC mode, and most **debug** commands take no arguments. For example, beginning in privileged EXEC mode, enter this command to enable the debugging for Switched Port Analyzer (SPAN):

```
Switch# debug span-session
```

The switch continues to generate output until you enter the **no** form of the command.

If you enable a **debug** command and no output appears, consider these possibilities:

- The switch might not be properly configured to generate the type of traffic you want to monitor. Use the **show running-config** command to check its configuration.
- Even if the switch is properly configured, it might not generate the type of traffic you want to monitor during the particular period that debugging is enabled. Depending on the feature you are debugging, you can use commands such as the TCP/IP **ping** command to generate network traffic.

To disable debugging of SPAN, enter this command in privileged EXEC mode:

```
Switch# no debug span-session
```

Alternately, in privileged EXEC mode, you can enter the **undebug** form of the command:

```
Switch# undebug span-session
```

To display the state of each debugging option, enter this command in privileged EXEC mode:

```
Switch# show debugging
```

Enabling All-System Diagnostics

Beginning in privileged EXEC mode, enter this command to enable all-system diagnostics:

```
Switch# debug all
```

**Caution**

Because debugging output takes priority over other network traffic, and because the **debug all** privileged EXEC command generates more output than any other **debug** command, it can severely diminish switch performance or even render it unusable. In virtually all cases, it is best to use more specific **debug** commands.

The **no debug all** privileged EXEC command disables all diagnostic output. Using the **no debug all** command is a convenient way to ensure that you have not accidentally left any **debug** commands enabled.

Redirecting Debug and Error Message Output

By default, the network server sends the output from **debug** commands and system error messages to the console. If you use this default, you can use a virtual terminal connection to monitor debug output instead of connecting to the console port.

Possible destinations include the console, virtual terminals, internal buffer, and UNIX hosts running a syslog server. The syslog format is compatible with 4.3 Berkeley Standard Distribution (BSD) UNIX and its derivatives.



Note

Be aware that the debugging destination you use affects system overhead. Logging messages to the console produces very high overhead, whereas logging messages to a virtual terminal produces less overhead. Logging messages to a syslog server produces even less, and logging to an internal buffer produces the least overhead of any method.

For more information about system message logging, see [Chapter 24, “Configuring System Message Logging.”](#)

Using the show platform forward Command

The output from the **show platform forward** privileged EXEC command provides some useful information about the forwarding results if a packet entering an interface is sent through the system. Depending upon the parameters entered about the packet, the output provides lookup table results and port maps used to calculate forwarding destinations, bitmaps, and egress information.



Note

For more syntax and usage information for the **show platform forward** command, refer to the switch command reference for this release.

Most of the information in the output from the command is useful mainly for technical support personnel, who have access to detailed information about the switch application-specific integrated circuits (ASICs). However, packet forwarding information can also be helpful in troubleshooting.

This is an example of the output from the **show platform forward** command on Gigabit Ethernet port 4 in VLAN 5 when the packet entering that port is addressed to unknown MAC addresses. The packet should be flooded to all other ports in VLAN 5.

```
Switch# show platform forward gigabitethernet0/4 vlan 5 1.1.1 2.2.2 ip 13.1.1.1 13.2.2.2
udp 10 20
Global Port Number:24, Asic Number:5
Src Real Vlan Id:5, Mapped Vlan Id:5

Ingress:
  Lookup          Key-Used          Index-Hit  A-Data
InptACL  40_0D020202_0D010101-00_40000014_000A0000    01FFA    03000000
L2Local  80_00050002_00020002-00_00000000_00000000    00C71    0000002B
Station Descriptor:02340000, DestIndex:0239, RewriteIndex:F005

=====
Egress:Asic 2, switch 1
Output Packets:

-----
Packet 1
```

```

      Lookup                Key-Used                Index-Hit  A-Data
OutputACL 50_0D020202_0D010101-00_40000014_000A0000      01FFE  03000000

Port      Vlan      SrcMac      DstMac      Cos  Dscp
Gi0/3      0005 0001.0001.0001  0002.0002.0002

-----

Packet 2
      Lookup                Key-Used                Index-Hit  A-Data
OutputACL 50_0D020202_0D010101-00_40000014_000A0000      01FFE  03000000

Port      Vlan      SrcMac      DstMac      Cos  Dscp
Gi0/4      0005 0001.0001.0001  0002.0002.0002

-----

Packet 3
      Lookup                Key-Used                Index-Hit  A-Data
OutputACL 50_0D020202_0D010101-00_40000014_000A0000      01FFE  03000000

Port      Vlan      SrcMac      DstMac      Cos  Dscp
Gi0/2      0005 0001.0001.0001  0002.0002.0002

-----
<output truncated>
-----

Packet 10
      Lookup                Key-Used                Index-Hit  A-Data
OutputACL 50_0D020202_0D010101-00_40000014_000A0000      01FFE  03000000
Packet dropped due to failed DEJA_VU Check on Gi0/4

```

This is an example of the output when the packet coming in on Gigabit Ethernet port 4 in VLAN 5 is sent to an address already learned on the VLAN on another port. It should be forwarded from the port on which the address was learned.

```
Switch# show platform forward gigabitethernet0/4 vlan 5 1.1.1 0009.43a8.0145 ip 13.1.1.1 13.2.2.2 udp 10 20
```

```
Global Port Number:24, Asic Number:5
```

```
Src Real Vlan Id:5, Mapped Vlan Id:5
```

```

Ingress:
      Lookup                Key-Used                Index-Hit  A-Data
InptACL  40_0D020202_0D010101-00_40000014_000A0000      01FFA  03000000
L2Local  80_00050009_43A80145-00_00000000_00000000      00086  02010197
Station Descriptor:F0050003, DestIndex:F005, RewriteIndex:0003

```

```
=====
```

```
Egress:Asic 3, switch 1
```

```
Output Packets:
```

```

-----

Packet 1
      Lookup                Key-Used                Index-Hit  A-Data
OutputACL 50_0D020202_0D010101-00_40000014_000A0000      01FFE  03000000

Port      Vlan      SrcMac      DstMac      Cos  Dscp
Gi0/3      0005 0001.0001.0001  0009.43A8.0145

```

Using the crashinfo File

The crashinfo file saves information that helps Cisco technical support representatives to debug problems that caused the Cisco IOS image to fail (crash). The switch writes the crash information to the console at the time of the failure, and the file is created the next time you boot the Cisco IOS image after the failure (instead of while the system is failing).

The information in the file includes the Cisco IOS image name and version that failed, a list of the processor registers, and a stack trace. You can provide this information to the Cisco technical support representative by using the **show tech-support** privileged EXEC command.

All crashinfo files are kept in this directory on the Flash file system:

flash:/crashinfo/crashinfo_*n* where *n* is a sequence number.

Each new crashinfo file that is created uses a sequence number that is larger than any previously existing sequence number, so the file with the largest sequence number describes the most recent failure. Version numbers are used instead of a timestamp because the switches do not include a real-time clock. You cannot change the name of the file that the system will use when it creates the file. However, after the file is created, you can use the **rename** privileged EXEC command to rename it, but the contents of the renamed file will not be displayed by the **show stacks** or the **show tech-support** privileged EXEC command. You can delete crashinfo files by using the **delete** privileged EXEC command.

You can display the most recent crashinfo file (that is, the file with the highest sequence number at the end of its filename) by entering the **show stacks** or the **show tech-support** privileged EXEC command. You also can access the file by using any command that can copy or display files, such as the **more** or the **copy** privileged EXEC command.



Supported MIBs

This appendix lists the supported management information base (MIBs) for this release on the Catalyst 2970 switch. It contains these sections:

- [MIB List, page A-1](#)
- [Using FTP to Access the MIB Files, page A-3](#)

MIB List

- BRIDGE-MIB (RFC1493)



Note The BRIDGE-MIB supports the context of a single VLAN. By default, SNMP messages using the configured community string always provide information for VLAN 1. To obtain the BRIDGE-MIB information for other VLANs, for example VLAN x, use this community string in the SNMP message: configured community string @x.

- CISCO-CDP-MIB
- CISCO-CLUSTER-MIB
- CISCO-CONFIG-COPY-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-ENTITY-FRU-CONTROL-MIB
- CISCO-ENVMON-MIB
- CISCO-FLASH-MIB (Flash memory on all switches is modeled as removable Flash memory.)
- CISCO-FTP-CLIENT-MIB
- CISCO-IGMP-FILTER-MIB
- CISCO-IMAGE-MIB
- CISCO IP-STAT-MIB
- CISCO-L2L3-INTERFACE-MIB
- CISCO-LACP-MIB
- CISCO-MAC-NOTIFICATION-MIB
- CISCO-MEMORY-POOL-MIB

- CISCO-PAGP-MIB
- CISCO-PING-MIB
- CISCO-PROCESS-MIB
- CISCO-RTTMON-MIB
- CISCO-STACKMAKER-MIB
- CISCO-STP-EXTENSIONS-MIB
- CISCO-SYSLOG-MIB
- CISCO-TCP-MIB
- CISCO-VLAN-IFTABLE-RELATIONSHIP-MIB
- CISCO-VLAN-MEMBERSHIP-MIB
- CISCO-VTP-MIB
- ENTITY-MIB
- ETHERLIKE_MIB
- IEEE8023-LACP-MIB
- IF-MIB (In and out counters for VLANs are not supported.)
- IGMP-MIB
- OLD-CISCO-CHASSIS-MIB
- OLD-CISCO-FLASH-MIB
- OLD-CISCO-INTERFACES-MIB
- OLD-CISCO-IP-MIB
- OLD-CISCO-SYS-MIB
- OLD-CISCO-TCP-MIB
- OLD-CISCO-TS-MIB
- RFC1213-MIB (Functionality is as per the agent capabilities specified in the CISCO-RFC1213-CAPABILITY.my.)
- RMON-MIB
- RMON2-MIB
- SNMP-FRAMEWORK-MIB
- SNMP-MPD-MIB
- SNMP-NOTIFICATION-MIB
- SNMP-TARGET-MIB
- SNMPv2-MIB
- TCP-MIB
- UDP-MIB

**Note**

You can also use this URL for a list of supported MIBs for the Catalyst 2970 switch:
<ftp://ftp.cisco.com/pub/mibs/supportlists/cat2970/cat2970-supportlist.html>

You can access other information about MIBs and Cisco products on the Cisco web site:
<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

Using FTP to Access the MIB Files

You can obtain each MIB file by using this procedure:

-
- | | |
|---------------|---|
| Step 1 | Use FTP to access the server ftp.cisco.com . |
| Step 2 | Log in with the username anonymous . |
| Step 3 | Enter your e-mail username when prompted for the password. |
| Step 4 | At the <code>ftp></code> prompt, change directories to /pub/mibs/v1 and /pub/mibs/v2 . |
| Step 5 | Use the get <i>MIB_filename</i> command to obtain a copy of the MIB file. |
-



Working with the Cisco IOS File System, Configuration Files, and Software Images

This appendix describes how to manipulate the Catalyst 2970 Flash file system, how to copy configuration files, and how to archive (upload and download) software images to a switch.



Note

For complete syntax and usage information for the commands used in this chapter, refer to the switch command reference for this release and the *Cisco IOS Configuration Fundamentals Command Reference for Release 12.1*.

This appendix consists of these sections:

- [Working with the Flash File System, page B-1](#)
- [Working with Configuration Files, page B-8](#)
- [Working with Software Images, page B-20](#)

Working with the Flash File System

The Flash file system is a single Flash device on which you can store files. It also provides several commands to help you manage software image and configuration files. The default Flash file system on the switch is named *flash:*.

This section contains this information:

- [Displaying Available File Systems, page B-2](#)
- [Setting the Default File System, page B-3](#)
- [Displaying Information about Files on a File System, page B-3](#)
- [Creating and Removing Directories, page B-4](#)
- [Copying Files, page B-4](#)
- [Deleting Files, page B-5](#)
- [Creating, Displaying, and Extracting tar Files, page B-5](#)
- [Displaying the Contents of a File, page B-7](#)

Displaying Available File Systems

To display the available file systems on your switch, use the **show file systems** privileged EXEC command as shown in this example.

```
Switch# show file systems
File Systems:
Size(b)      Free(b)      Type  Flags  Prefixes
*  15998976   5135872      flash  rw    flash:
          -         -      opaque  rw    bs:
          -         -      opaque  rw    vb:
          524288   520138      nvram   rw    nvram:
          -         -      network  rw    tftp:
          -         -      opaque  rw    null:
          -         -      opaque  rw    system:
          -         -      opaque  ro    xmodem:
          -         -      opaque  ro    ymodem:
```

Table B-1 *show file systems Field Descriptions*

Field	Value
Size(b)	Amount of memory in the file system in bytes.
Free(b)	Amount of free memory in the file system in bytes.
Type	Type of file system. flash —The file system is for a Flash memory device. nvram —The file system is for a nonvolatile RAM (NVRAM) device. opaque —The file system is a locally generated <i>pseudo</i> file system (for example, the <i>system</i>) or a download interface, such as brimux. unknown —The file system is an unknown type.
Flags	Permission for file system. ro —read-only. rw —read/write. wo —write-only.
Prefixes	Alias for file system. flash: —Flash file system. nvram: —NVRAM. null: —Null destination for copies. You can copy a remote file to null to determine its size. rcp: —Remote Copy Protocol (RCP) network server. system: —Contains the system memory, including the running configuration. tftp: —Trivial File Transfer Protocol (TFTP) network server. xmodem: —Obtain the file from a network machine by using the XMODEM protocol. ymodem: —Obtain the file from a network machine by using the YMODEM protocol.

Setting the Default File System

You can specify the file system or directory that the system uses as the default file system by using the **cd filesystem:** privileged EXEC command. You can set the default file system to omit the *filesystem:* argument from related commands. For example, for all privileged EXEC commands that have the optional *filesystem:* argument, the system uses the file system specified by the **cd** command.

By default, the default file system is *flash:*.

You can display the current default file system as specified by the **cd** command by using the **pwd** privileged EXEC command.

Displaying Information about Files on a File System

You can view a list of the contents of a file system before manipulating its contents. For example, before copying a new configuration file to Flash memory, you might want to verify that the file system does not already contain a configuration file with the same name. Similarly, before copying a Flash configuration file to another location, you might want to verify its filename for use in another command.

To display information about files on a file system, use one of the privileged EXEC commands in [Table B-2](#):

Table B-2 Commands for Displaying Information About Files

Command	Description
dir [/all] [filesystem:][filename]	Display a list of files on a file system.
show file systems	Display more information about each of the files on a file system.
show file information file-url	Display information about a specific file.
show file descriptors	Display a list of open file descriptors. File descriptors are the internal representations of open files. You can use this command to see if another user has a file open.

Changing Directories and Displaying the Working Directory

Beginning in privileged EXEC mode, follow these steps to change directories and display the working directory.

	Command	Purpose
Step 1	dir filesystem:	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board Flash device.
Step 2	cd new_configs	Change to the directory of interest. The command example shows how to change to the directory named <i>new_configs</i> .
Step 3	pwd	Display the working directory.

Creating and Removing Directories

Beginning in privileged EXEC mode, follow these steps to create and remove a directory:

	Command	Purpose
Step 1	dir <i>filesystem:</i>	Display the directories on the specified file system. For <i>filesystem:</i> , use flash: for the system board Flash device.
Step 2	mkdir <i>old_configs</i>	Create a new directory. The command example shows how to create the directory named <i>old_configs</i> . Directory names are case sensitive. Directory names are limited to 45 characters between the slashes (/); the name cannot contain control characters, spaces, deletes, slashes, quotes, semicolons, or colons.
Step 3	dir <i>filesystem:</i>	Verify your entry.

To delete a directory with all its files and subdirectories, use the **delete /force /recursive** *filesystem:/file-url* privileged EXEC command.

Use the **/recursive** keyword to delete the named directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the name of the directory to be deleted. All the files in the directory and the directory are removed.



Caution

When files and directories are deleted, their contents cannot be recovered.

Copying Files

To copy a file from a source to a destination, use the **copy** *source-url destination-url* privileged EXEC command. For the source and destination URLs, you can use **running-config** and **startup-config** keyword shortcuts. For example, the **copy running-config startup-config** command saves the currently running configuration file to the NVRAM section of Flash memory to be used as the configuration during system initialization.

You can also copy from special file systems (**xmodem:**, **ymodem:**) as the source for the file from a network machine that uses the XMODEM or YMODEM protocol.

Network file system URLs include **ftp:**, **rcp:**, and **tftp:** and have these syntaxes:

- File Transfer Protocol (FTP)—**ftp:**[[//username [:password]@location]/directory]/filename
- Remote Copy Protocol (RCP)—**rcp:**[[//username@location]/directory]/filename
- Trivial File Transfer Protocol (TFTP)—**tftp:**[[//location]/directory]/filename

Local writable file systems include **flash:**.

Some invalid combinations of source and destination exist. Specifically, you cannot copy these combinations:

- From a running configuration to a running configuration
- From a startup configuration to a startup configuration
- From a device to the same device (for example, the **copy flash: flash:** command is invalid)

For specific examples of using the **copy** command with configuration files, see the [“Working with Configuration Files” section on page B-8](#).

To copy software images either by downloading a new version or uploading the existing one, use the **archive download-sw** or the **archive upload-sw** privileged EXEC command. For more information, see the [“Working with Software Images” section on page B-20](#).

Deleting Files

When you no longer need a file on a Flash memory device, you can permanently delete it. To delete a file or directory from a specified Flash device, use the **delete** [**/force**] [**/recursive**] [*filesystem:*]/*file-url* privileged EXEC command.

Use the **/recursive** keyword for deleting a directory and all subdirectories and the files contained in it. Use the **/force** keyword to suppress the prompting that confirms a deletion of each file in the directory. You are prompted only once at the beginning of this deletion process. Use the **/force** and **/recursive** keywords for deleting old software images that were installed by using the **archive download-sw** command but are no longer needed.

If you omit the *filesystem:* option, the switch uses the default device specified by the **cd** command. For *file-url*, you specify the path (directory) and the name of the file to be deleted.

When you attempt to delete any files, the system prompts you to confirm the deletion.



Caution

When files are deleted, their contents cannot be recovered.

This example shows how to delete the file *myconfig* from the default Flash memory device:

```
Switch# delete myconfig
```

Creating, Displaying, and Extracting tar Files

You can create a tar file and write files into it, list the files in a tar file, and extract the files from a tar file as described in the next sections.



Note

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

Creating a tar File

To create a tar file and write files into it, use this privileged EXEC command:

archive tar /create destination-url flash:/file-url

For *destination-url*, specify the destination URL alias for the local or network file system and the name of the tar file to create. These options are supported:

- For the local Flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:[[/username[:password]@location]/directory]/tar-filename.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rctp:[[/username@location]/directory]/tar-filename.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:[[/location]/directory]/tar-filename.tar**

The *tar-filename.tar* is the tar file to be created.

For **flash:/file-url**, specify the location on the local Flash file system from which the new tar file is created. You can also specify an optional list of files or directories within the source directory to write to the new tar file. If none are specified, all files and directories at this level are written to the newly created tar file.

This example shows how to create a tar file. This command writes the contents of the *new-configs* directory on the local Flash device to a file named *saved.tar* on the TFTP server at 172.20.10.30:

```
Switch# archive tar /create tftp:172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a tar File

To display the contents of a tar file on the screen, use this privileged EXEC command:

archive tar /table source-url

For *source-url*, specify the source URL alias for the local or network file system. These options are supported:

- For the local Flash file system, the syntax is **flash:**
- For the File Transfer Protocol (FTP), the syntax is **ftp:[[/username[:password]@location]/directory]/tar-filename.tar**
- For the Remote Copy Protocol (RCP), the syntax is **rctp:[[/username@location]/directory]/tar-filename.tar**
- For the Trivial File Transfer Protocol (TFTP), the syntax is **tftp:[[/location]/directory]/tar-filename.tar**

The *tar-filename.tar* is the tar file to display.

You can also limit the display of the files by specifying an optional list of files or directories after the tar file; then only those files appear. If none are specified, all files and directories appear.

This example shows how to display the contents of a switch tar file that is in Flash memory:

```
Switch# archive tar /table flash:c2970-i612-mz.121-6.AX1.tar
```

```
info (219 bytes)
c2970-i612-mz.121-6.AX1/ (directory)
```

```
c2970-i612-mz.121-6.AX1/html/ (directory)
c2970-i612-mz.121-6.AX1/html/foo.html (0 bytes)
c2970-i612-mz.121-6.AX1/c2970-i612-mz.121-6.AX1.bin (610856 bytes)
c2970-i612-mz.121-6.AX1/info (219 bytes)
```

This example shows how to display only the /html directory and its contents:

```
Switch# archive tar /table flash:c2970-tv0-m.tar c2970-i612-mz.121-6.AX1/html
c2970-i612-mz.121-6.AX1/html/ (directory)
c2970-i612-mz.121-6.AX1/html/foo.html (0 bytes)
```

Extracting a tar File

To extract a tar file into a directory on the Flash file system, use this privileged EXEC command:

archive tar /xtract *source-url* **flash:***file-url* [*dir/file...*]

For *source-url*, specify the source URL alias for the local file system. These options are supported:

- For the local Flash file system, the syntax is
flash:
- For the File Transfer Protocol (FTP), the syntax is
ftp:*[/username[:password]@location]/directory/tar-filename.tar*
- For the Remote Copy Protocol (RCP), the syntax is
rcp:*[/username@location]/directory/tar-filename.tar*
- For the Trivial File Transfer Protocol (TFTP), the syntax is
tftp:*[/location]/directory/tar-filename.tar*

The *tar-filename.tar* is the tar file from which to extract files.

For **flash:***file-url* [*dir/file...*], specify the location on the local Flash file system into which the tar file is extracted. Use the *dir/file...* option to specify an optional list of files or directories within the tar file to be extracted. If none are specified, all files and directories are extracted.

This example shows how to extract the contents of a tar file located on the TFTP server at 172.20.10.30. This command extracts just the *new-configs* directory into the root directory on the local Flash file system. The remaining files in the *saved.tar* file are ignored.

```
Switch# archive tar /xtract tftp://172.20.10.30/saved.tar flash:/new-configs
```

Displaying the Contents of a File

To display the contents of any readable file, including a file on a remote file system, use the **more** [*/ascii* | */binary* | */ebcdic*] *file-url* privileged EXEC command:

This example shows how to display the contents of a configuration file on a TFTP server:

```
Switch# more tftp://serverA/hampton/savedconfig
!
! Saved configuration on server
!
version 11.3
service timestamps log datetime localtime
service linenummer
service udp-small-servers
service pt-vty-logging
!

<output truncated>
```

Working with Configuration Files

This section describes how to create, load, and maintain configuration files.

Configuration files contain commands entered to customize the function of the Cisco IOS software. A way to create a basic configuration file is to use the **setup** program or to enter the **setup** privileged EXEC command. For more information, see [Chapter 4, “Assigning the Switch IP Address and Default Gateway.”](#)

You can copy (*download*) configuration files from a TFTP, FTP, or RCP server to the running configuration or startup configuration of the switch. You might want to perform this for one of these reasons:

- To restore a backed-up configuration file.
- To use the configuration file for another switch. For example, you might add another switch to your network and want it to have a configuration similar to the original switch. By copying the file to the new switch, you can change the relevant parts rather than recreating the whole file.
- To load the same configuration commands on all the switches in your network so that all the switches have similar configurations.

You can copy (*upload*) configuration files from the switch to a file server by using TFTP, FTP, or RCP. You might perform this task to back up a current configuration file to a server before changing its contents so that you can later restore the original configuration file from the server.

The protocol you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- [Guidelines for Creating and Using Configuration Files, page B-9](#)
- [Configuration File Types and Location, page B-9](#)
- [Creating a Configuration File By Using a Text Editor, page B-10](#)
- [Copying Configuration Files By Using TFTP, page B-10](#)
- [Copying Configuration Files By Using FTP, page B-12](#)
- [Copying Configuration Files By Using RCP, page B-16](#)
- [Clearing Configuration Information, page B-19](#)

Guidelines for Creating and Using Configuration Files

Creating configuration files can aid in your switch configuration. Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration.

Use these guidelines when creating a configuration file:

- We recommend that you connect through the console port for the initial configuration of the switch. If you are accessing the switch through a network connection instead of through a direct connection to the console port, keep in mind that some configuration changes (such as changing the switch IP address or disabling ports) can cause a loss of connectivity to the switch.
- If no password has been set on the switch, we recommend that you set one by using the **enable secret** *secret-password* global configuration command.



Note

The **copy {ftp: | rcp: | tftp:} system:running-config** privileged EXEC command loads the configuration files on the switch as if you were entering the commands at the command line. The switch does not erase the existing running configuration before adding the commands. If a command in the copied configuration file replaces a command in the existing configuration file, the existing command is erased. For example, if the copied configuration file contains a different IP address in a particular command than the existing configuration, the IP address in the copied configuration is used. However, some commands in the existing configuration might not be replaced or negated. In this case, the resulting configuration file is a mixture of the existing configuration file and the copied configuration file, with the copied configuration file having precedence.

To restore a configuration file to an exact copy of a file stored on a server, copy the configuration file directly to the startup configuration (by using the **copy {ftp: | rcp: | tftp:} nvram:startup-config** privileged EXEC command), and reload the switch.

Configuration File Types and Location

Startup configuration files are used during system startup to configure the software. Running configuration files contain the current configuration of the software. The two configuration files can be different. For example, you might want to change the configuration for a short time period rather than permanently. In this case, you would change the running configuration but not save the configuration by using the **copy running-config startup-config** privileged EXEC command.

The running configuration is saved in DRAM; the startup configuration is stored in the NVRAM section of Flash memory.

Creating a Configuration File By Using a Text Editor

When creating a configuration file, you must list commands logically so that the system can respond appropriately. This is one method of creating a configuration file:

-
- | | |
|--------|--|
| Step 1 | Copy an existing configuration from a switch to a server.

For more information, see the “Downloading the Configuration File By Using TFTP” section on page B-11, the “Downloading a Configuration File By Using FTP” section on page B-13, or the “Downloading a Configuration File By Using RCP” section on page B-17. |
| Step 2 | Open the configuration file in a text editor, such as vi or emacs on UNIX or Notepad on a PC. |
| Step 3 | Extract the portion of the configuration file with the desired commands, and save it in a new file. |
| Step 4 | Copy the configuration file to the appropriate server location. For example, copy the file to the TFTP directory on the workstation (usually /tftpboot on a UNIX workstation). |
| Step 5 | Make sure the permissions on the file are set to world-read. |
-

Copying Configuration Files By Using TFTP

You can configure the switch by using configuration files you create, download from another switch, or download from a TFTP server. You can copy (upload) configuration files to a TFTP server for storage.

This section includes this information:

- [Preparing to Download or Upload a Configuration File By Using TFTP, page B-10](#)
- [Downloading the Configuration File By Using TFTP, page B-11](#)
- [Uploading the Configuration File By Using TFTP, page B-11](#)

Preparing to Download or Upload a Configuration File By Using TFTP

Before you begin downloading or uploading a configuration file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the /etc/inetd.conf file contains this line:

```
tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
```

Make sure that the /etc/services file contains this line:

```
tftp 69/udp
```



Note

You must restart the inetd daemon after modifying the /etc/inetd.conf and /etc/services files. To restart the daemon, either stop the inetd process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.

- Ensure that the configuration file to be downloaded is in the correct directory on the TFTP server (usually /tftpboot on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.
- Before uploading the configuration file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch filename** command, where *filename* is the name of the file you will use when uploading it to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading the Configuration File By Using TFTP

To configure the switch by using a configuration file downloaded from a TFTP server, follow these steps:

-
- Step 1** Copy the configuration file to the appropriate TFTP directory on the workstation.
- Step 2** Verify that the TFTP server is properly configured by referring to the [“Preparing to Download or Upload a Configuration File By Using TFTP” section on page B-10](#).
- Step 3** Log into the switch through the console port or a Telnet session.
- Step 4** Download the configuration file from the TFTP server to configure the switch.

Specify the IP address or host name of the TFTP server and the name of the file to download.

Use one of these privileged EXEC commands:

- **copy tftp:[*[[//location]/directory]/filename*] system:running-config**
- **copy tftp:[*[[//location]/directory]/filename*] nvram:startup-config**

The configuration file downloads, and the commands are executed as the file is parsed line-by-line.

This example shows how to configure the software from the file *tokyo-config* at IP address 172.16.2.155:

```
Switch# copy tftp://172.16.2.155/tokyo-config system:running-config
Configure using tokyo-config from 172.16.2.155? [confirm] y
Booting tokyo-config from 172.16.2.155:!!! [OK - 874/16000 bytes]
```

Uploading the Configuration File By Using TFTP

To upload a configuration file from a switch to a TFTP server for storage, follow these steps:

-
- Step 1** Verify that the TFTP server is properly configured by referring to the [“Preparing to Download or Upload a Configuration File By Using TFTP” section on page B-10](#).
- Step 2** Log into the switch through the console port or a Telnet session.

- Step 3** Upload the switch configuration to the TFTP server. Specify the IP address or host name of the TFTP server and the destination filename.

Use one of these privileged EXEC commands:

- **copy system:running-config tftp:**[[*location*]/*directory*]/*filename*
- **copy nvram:startup-config tftp:**[[*location*]/*directory*]/*filename*

The file is uploaded to the TFTP server.

This example shows how to upload a configuration file from a switch to a TFTP server:

```
Switch# copy system:running-config tftp://172.16.2.155/tokyo-config
Write file tokyo-config on host 172.16.2.155? [confirm] y
#
Writing tokyo-config!!! [OK]
```

Copying Configuration Files By Using FTP

You can copy configuration files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy a configuration file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **copy** command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept your FTP write request.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **copy** command if you want to specify only a username for that copy operation.

If the server has a directory structure, the configuration file is written to or copied from the directory associated with the username on the server. For example, if the configuration file resides in the home directory of a user on the server, specify that user's name as the remote username.

For more information, refer to the documentation for your FTP server.

This section includes this information:

- [Preparing to Download or Upload a Configuration File By Using FTP, page B-13](#)
- [Downloading a Configuration File By Using FTP, page B-13](#)
- [Uploading a Configuration File By Using FTP, page B-15](#)

Preparing to Download or Upload a Configuration File By Using FTP

Before you begin downloading or uploading a configuration file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username username** global configuration command during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.
- When you upload a configuration file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using FTP:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using FTP” section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode on the switch. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.

	Command	Purpose
Step 6	end	Return to privileged EXEC mode.
Step 7	copy ftp:[username[:password]@]location/]directory /filename] system:running-config or copy ftp:[username[:password]@]location/]directory /filename] nvram:startup-config	Using FTP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and to load and run those commands on the switch:

```
Switch# copy ftp://netadmin1:mypass@172.16.101.101/host1-config system:running-config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:![OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by ftp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. The software copies the configuration file *host2-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the switch startup configuration.

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin1
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy ftp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-config]? host2-config
Configure using host2-config from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-config:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-config by ftp from
172.16.101.101
```

Uploading a Configuration File By Using FTP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using FTP:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using FTP” section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 6	end	Return to privileged EXEC mode.
Step 7	copy system:running-config ftp:[[/[username[:password]@]location]/directory] /filename] or copy nvram:startup-config ftp:[[/[username[:password]@]location]/directory] /filename]	Using FTP, store the switch running or startup configuration file to the specified location.

This example shows how to copy the running configuration file named *switch2-config* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config ftp://netadmin1:mypass@172.16.101.101/switch2-config
Write file switch2-config on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server by using FTP to copy the file:

```
Switch# configure terminal
Switch(config)# ip ftp username netadmin2
Switch(config)# ip ftp password mypass
Switch(config)# end
Switch# copy nvram:startup-config ftp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-config]?
Write file switch2-config on host 172.16.101.101?[confirm]
![OK]
```

Copying Configuration Files By Using RCP

The Remote Copy Protocol (RCP) provides another method of downloading, uploading, and copying configuration files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

The RCP requires a client to send a remote username with each RCP request to a server. When you copy a configuration file from the switch to a server, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **copy** command if a username is specified.
- The username set by the **ip rcmd remote-username username** global configuration command if the command is configured.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For a successful RCP copy request, you must define an account on the network server for the remote username. If the server has a directory structure, the configuration file is written to or copied from the directory associated with the remote username on the server. For example, if the configuration file is in the home directory of a user on the server, specify that user's name as the remote username.

This section includes this information:

- [Preparing to Download or Upload a Configuration File By Using RCP, page B-16](#)
- [Downloading a Configuration File By Using RCP, page B-17](#)
- [Uploading a Configuration File By Using RCP, page B-18](#)

Preparing to Download or Upload a Configuration File By Using RCP

Before you begin downloading or uploading a configuration file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username username** global configuration command to be used during all copy operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the RCP username. Include the username in the **copy** command if you want to specify a username for only that copy operation.

- When you upload a file to the RCP server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the `.rhosts` file for the remote user on the RCP server. For example, suppose that the switch contains these configuration lines:

```
hostname Switch1
ip rcmd remote-username User0
```

If the switch IP address translates to *Switch1.company.com*, the `.rhosts` file for User0 on the RCP server should contain this line:

```
Switch1.company.com Switch1
```

For more information, refer to the documentation for your RCP server.

Downloading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to download a configuration file by using RCP:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using RCP” section on page B-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy rcp:[[//[<i>username@</i>]<i>location</i>]/<i>directory</i>]/<i>filename</i>] system:running-config or copy rcp:[[//[<i>username@</i>]<i>location</i>]/<i>directory</i>]/<i>filename</i>] nvrn:startup-config	Using RCP, copy the configuration file from a network server to the running configuration or to the startup configuration file.

This example shows how to copy a configuration file named *host1-config* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 and load and run those commands on the switch:

```
Switch# copy rcp://netadmin1@172.16.101.101/host1-config system:running-config
Configure using host1-config from 172.16.101.101? [confirm]
Connected to 172.16.101.101
Loading 1112 byte file host1-config:[OK]
Switch#
%SYS-5-CONFIG: Configured from host1-config by rcp from 172.16.101.101
```

This example shows how to specify a remote username of *netadmin1*. Then it copies the configuration file *host2-conf* from the *netadmin1* directory on the remote server with an IP address of 172.16.101.101 to the startup configuration:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin1
Switch(config)# end
Switch# copy rcp: nvram:startup-config
Address of remote host [255.255.255.255]? 172.16.101.101
Name of configuration file[rtr2-conf]? host2-conf
Configure using host2-conf from 172.16.101.101?[confirm]
Connected to 172.16.101.101
Loading 1112 byte file host2-conf:![OK]
[OK]
Switch#
%SYS-5-CONFIG_NV:Non-volatile store configured from host2-conf by rcp from
172.16.101.101
```

Uploading a Configuration File By Using RCP

Beginning in privileged EXEC mode, follow these steps to upload a configuration file by using RCP:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using RCP” section on page B-16.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy system:running-config rcp:[[/[username@]location]/directory]/filename] or copy nvram:startup-config rcp:[[/[username@]location]/directory]/filename]	Using RCP, copy the configuration file from a switch running or startup configuration file to a network server.

This example shows how to copy the running configuration file named *switch2-conf* to the *netadmin1* directory on the remote host with an IP address of 172.16.101.101:

```
Switch# copy system:running-config rcp://netadmin1@172.16.101.101/switch2-conf
Write file switch-conf on host 172.16.101.101?[confirm]
Building configuration...[OK]
Connected to 172.16.101.101
Switch#
```

This example shows how to store a startup configuration file on a server:

```
Switch# configure terminal
Switch(config)# ip rcmd remote-username netadmin2
Switch(config)# end
Switch# copy nvram:startup-config rcp:
Remote host[]? 172.16.101.101
Name of configuration file to write [switch2-config]?
Write file switch2-config on host 172.16.101.101?[confirm]
![OK]
```

Clearing Configuration Information

You can clear the configuration information from the startup configuration. If you reboot the switch with no startup configuration, the switch enters the setup program so that you can reconfigure the switch with all new settings.

Clearing the Startup Configuration File

To clear the contents of your startup configuration, use the **erase nvram:** or the **erase startup-config** privileged EXEC command.



Caution

You cannot restore the startup configuration file after it has been deleted.

Deleting a Stored Configuration File

To delete a saved configuration from Flash memory, use the **delete flash:filename** privileged EXEC command. Depending on the setting of the **file prompt** global configuration command, you might be prompted for confirmation before you delete a file. By default, the switch prompts for confirmation on destructive file operations. For more information about the **file prompt** command, refer to the *Cisco IOS Command Reference for Release 12.1*.



Caution

You cannot restore a file after it has been deleted.

Working with Software Images

This section describes how to archive (download and upload) software image files, which contain the system software, Cisco IOS code, and the web management HTML files.

**Note**

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

You download a switch image file from a TFTP, FTP, or RCP server to upgrade the switch software. You can replace the current image with the new one or keep the current image in Flash memory after a download.

You upload a switch image file to a TFTP, FTP, or RCP server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

The protocol you use depends on which type of server you are using. The FTP and RCP transport mechanisms provide faster performance and more reliable delivery of data than TFTP. These improvements are possible because FTP and RCP are built on and use the Transmission Control Protocol/Internet Protocol (TCP/IP) stack, which is connection-oriented.

This section includes this information:

- [Image Location on the Switch, page B-20](#)
- [tar File Format of Images on a Server or Cisco.com, page B-21](#)
- [Copying Image Files By Using TFTP, page B-22](#)
- [Copying Image Files By Using FTP, page B-25](#)
- [Copying Image Files By Using RCP, page B-29](#)

**Note**

For a list of software images and the supported upgrade paths, refer to the release notes that shipped with your switch.

Image Location on the Switch

The Cisco IOS image is stored as a *.bin* file in a directory that shows the version number. A subdirectory contains the HTML files needed for web management. The image is stored on the system board Flash memory (flash:).

You can use the **show version** privileged EXEC command to see the software version that is currently running on your switch. In the display, check the line that begins with `System image file is...`. It shows the directory name in Flash memory where the image is stored.

You can also use the **dir filesystem:** privileged EXEC command to see the directory names of other software images you might have stored in Flash memory.

tar File Format of Images on a Server or Cisco.com

Software images located on a server or downloaded from Cisco.com are provided in a tar file format, which contains these files:

- An *info* file, which serves as a table of contents for the tar file
- One or more subdirectories containing other images and files, such as Cisco IOS images and web management files

This example shows some of the information contained in the info file. [Table B-3](#) provides additional details about this information:

```
system_type:0x00000000:c2970-i612-mz.121.11-AX
  image_family:C2970
  stacking_number:1.0
  info_end:
version_suffix:i612-121.11-AX
  version_directory:c2970-i612-mz.121.11-AX
  image_system_type_id:0x00000000
  image_name:c2970-i612-mz.121.11-AX.bin
  ios_image_file_size:2939392
  total_image_file_size:4884992
  image_feature:LAYER_2|MIN_DRAM_MEG=32
  image_family:C2970
  stacking_number:1.0
  board_ids:0x00000008
  info_end:
```



Note

Disregard the stacking_number field. It does not apply to the switch.

Table B-3 info File Description

Field	Description
version_suffix	Specifies the Cisco IOS image version string suffix
version_directory	Specifies the directory where the Cisco IOS image and the HTML subdirectory are installed
image_name	Specifies the name of the Cisco IOS image within the tar file
ios_image_file_size	Specifies the Cisco IOS image size in the tar file, which is an approximate measure of how much Flash memory is required to hold just the Cisco IOS image
total_image_file_size	Specifies the size of all the images (the Cisco IOS image and the HTML files) in the tar file, which is an approximate measure of how much Flash memory is required to hold them
image_feature	Describes the core functionality of the image
image_min_dram	Specifies the minimum amount of DRAM needed to run this image
image_family	Describes the family of products on which the software can be installed

Copying Image Files By Using TFTP

You can download a switch image from a TFTP server or upload the image from the switch to a TFTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes; this uploaded image can be used for future downloads to the same or another switch of the same type.



Note

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using TFTP, page B-22](#)
- [Downloading an Image File By Using TFTP, page B-23](#)
- [Uploading an Image File By Using TFTP, page B-24](#)

Preparing to Download or Upload an Image File By Using TFTP

Before you begin downloading or uploading an image file by using TFTP, do these tasks:

- Ensure that the workstation acting as the TFTP server is properly configured. On a Sun workstation, make sure that the `/etc/inetd.conf` file contains this line:

```
tftp dgram udp wait root /usr/etc/in.tftpd in.tftpd -p -s /tftpboot
```

Make sure that the `/etc/services` file contains this line:

```
tftp 69/udp
```



Note

You must restart the `inetd` daemon after modifying the `/etc/inetd.conf` and `/etc/services` files. To restart the daemon, either stop the `inetd` process and restart it, or enter a **fastboot** command (on the SunOS 4.x) or a **reboot** command (on Solaris 2.x or SunOS 5.x). For more information on the TFTP daemon, refer to the documentation for your workstation.

- Ensure that the switch has a route to the TFTP server. The switch and the TFTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the TFTP server by using the **ping** command.
- Ensure that the image to be downloaded is in the correct directory on the TFTP server (usually `/tftpboot` on a UNIX workstation).
- For download operations, ensure that the permissions on the file are set correctly. The permission on the file should be world-read.

- Before uploading the image file, you might need to create an empty file on the TFTP server. To create an empty file, enter the **touch filename** command, where *filename* is the name of the file you will use when uploading the image to the server.
- During upload operations, if you are overwriting an existing file (including an empty file, if you had to create one) on the server, ensure that the permissions on the file are set correctly. Permissions on the file should be world-write.

Downloading an Image File By Using TFTP

You can download a new image file and replace the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 3 to download a new image from a TFTP server and overwrite the existing image. To keep the current image, skip Step 3.

	Command	Purpose
Step 1		Copy the image to the appropriate TFTP directory on the workstation. Make sure the TFTP server is properly configured; see the “Preparing to Download or Upload an Image File By Using TFTP” section on page B-22.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	archive download-sw /overwrite /reload tftp:[//location]/directory/image-name.tar	Download the image file from the TFTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in Flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>//location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 4	archive download-sw /leave-old-sw /reload tftp:[//location]/directory/image-name.tar	Download the image file from the TFTP server to the switch, and keep the current image. <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>//location</i>, specify the IP address of the TFTP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

**Note**

If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the current running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image on the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive filesystem:/file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old image. All the files in the directory and the directory are removed.

**Caution**

For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using TFTP

You can upload an image from the switch to a TFTP server. You can later download this image to the switch or to another switch of the same type.

The upload feature should be used only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to a TFTP server:

	Command	Purpose
Step 1		Make sure the TFTP server is properly configured; see the “Preparing to Download or Upload an Image File By Using TFTP” section on page B-22.
Step 1		Log into the switch through the console port or a Telnet session.
Step 2	archive upload-sw tftp:[[/location]/directory]/image-name.tar	Upload the currently running switch image to the TFTP server. <ul style="list-style-type: none"> For <i>//location</i>, specify the IP address of the TFTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the HTML files. After these files are uploaded, the upload algorithm creates the tar file format.

**Caution**

For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using FTP

You can download a switch image from an FTP server or upload the image from the switch to an FTP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the switch or another switch of the same type.



Note

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using FTP, page B-25](#)
- [Downloading an Image File By Using FTP, page B-26](#)
- [Uploading an Image File By Using FTP, page B-28](#)

Preparing to Download or Upload an Image File By Using FTP

You can copy images files to or from an FTP server.

The FTP protocol requires a client to send a remote username and password on each FTP request to a server. When you copy an image file from the switch to a server by using FTP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip ftp username** *username* global configuration command if the command is configured.
- Anonymous.

The switch sends the first valid password in this list:

- The password specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a password is specified.
- The password set by the **ip ftp password** *password* global configuration command if the command is configured.
- The switch forms a password named *username@switchname.domain*. The variable *username* is the username associated with the current session, *switchname* is the configured host name, and *domain* is the domain of the switch.

The username and password must be associated with an account on the FTP server. If you are writing to the server, the FTP server must be properly configured to accept the FTP write request from you.

Use the **ip ftp username** and **ip ftp password** commands to specify a username and password for all copies. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.

If the server has a directory structure, the image file is written to or copied from the directory associated with the username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using FTP, do these tasks:

- Ensure that the switch has a route to the FTP server. The switch and the FTP server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the FTP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current FTP username is the one that you want to use for the FTP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new FTP username by using the **ip ftp username *username*** global configuration command. This new name will be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and you do not need to set the FTP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username for that operation only.
- When you upload an image file to the FTP server, it must be properly configured to accept the write request from the user on the switch.

For more information, refer to the documentation for your FTP server.

Downloading an Image File By Using FTP

You can download a new image file and overwrite the current image or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 7 to download a new image from an FTP server and overwrite the existing image. To keep the current image, skip Step 7.

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using FTP” section on page B-25.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
Step 7	archive download-sw /overwrite /reload ftp:[[/username[:password]@location]/directory] /image-name.tar	Download the image file from the FTP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in Flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username[:password], specify the username and password; these must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page B-25. • For @location, specify the IP address of the FTP server. • For directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.
Step 8	archive download-sw /leave-old-sw /reload ftp:[[/username[:password]@location]/directory] /image-name.tar	Download the image file from the FTP server to the switch, and keep the current image. <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username[:password], specify the username and password. These must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page B-25. • For @location, specify the IP address of the FTP server. • For directory/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device, whether or not it is the same as the new one, downloads the new image, and then reloads the software.



Note

If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough space to install the new image and keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old image during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive filesystem:/file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.



Caution

For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using FTP

You can upload an image from the switch to an FTP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature should be used only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an FTP server:

	Command	Purpose
Step 1		Verify that the FTP server is properly configured by referring to the “Preparing to Download or Upload a Configuration File By Using FTP” section on page B-13.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username or password (see Steps 4, 5, and 6).
Step 4	ip ftp username <i>username</i>	(Optional) Change the default remote username.
Step 5	ip ftp password <i>password</i>	(Optional) Change the default password.
Step 6	end	Return to privileged EXEC mode.
Step 7	archive upload-sw ftp:[[/[username[:password]@]location]/directory]/ image-name.tar	Upload the currently running switch image to the FTP server. <ul style="list-style-type: none"> For <i>//username:password</i>, specify the username and password. These must be associated with an account on the FTP server. For more information, see the “Preparing to Download or Upload an Image File By Using FTP” section on page B-25. For <i>@location</i>, specify the IP address of the FTP server. For <i>/directory/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of the software image to be stored on the server.

The **archive upload-sw** command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the HTML files. After these files are uploaded, the upload algorithm creates the tar file format.

**Caution**

For the download and upload algorithms to operate properly, do *not* rename image names.

Copying Image Files By Using RCP

You can download a switch image from an RCP server or upload the image from the switch to an RCP server.

You download a switch image file from a server to upgrade the switch software. You can overwrite the current image with the new one or keep the current image after a download.

You upload a switch image file to a server for backup purposes. You can use this uploaded image for future downloads to the same switch or another of the same type.

**Note**

Instead of using the **copy** privileged EXEC command or the **archive tar** privileged EXEC command, we recommend using the **archive download-sw** and **archive upload-sw** privileged EXEC commands to download and upload software image files.

This section includes this information:

- [Preparing to Download or Upload an Image File By Using RCP, page B-29](#)
- [Downloading an Image File By Using RCP, page B-31](#)
- [Uploading an Image File By Using RCP, page B-33](#)

Preparing to Download or Upload an Image File By Using RCP

RCP provides another method of downloading and uploading image files between remote hosts and the switch. Unlike TFTP, which uses User Datagram Protocol (UDP), a connectionless protocol, RCP uses TCP, which is connection-oriented.

To use RCP to copy files, the server from or to which you will be copying files must support RCP. The RCP copy commands rely on the rsh server (or daemon) on the remote system. To copy files by using RCP, you do not need to create a server for file distribution as you do with TFTP. You only need to have access to a server that supports the remote shell (rsh). (Most UNIX systems support rsh.) Because you are copying a file from one place to another, you must have read permission on the source file and write permission on the destination file. If the destination file does not exist, RCP creates it for you.

RCP requires a client to send a remote username on each RCP request to a server. When you copy an image from the switch to a server by using RCP, the Cisco IOS software sends the first valid username in this list:

- The username specified in the **archive download-sw** or **archive upload-sw** privileged EXEC command if a username is specified.
- The username set by the **ip rcmd remote-username username** global configuration command if the command is entered.
- The remote username associated with the current TTY (terminal) process. For example, if the user is connected to the router through Telnet and was authenticated through the **username** command, the switch software sends the Telnet username as the remote username.
- The switch host name.

For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. If the server has a directory structure, the image file is written to or copied from the directory associated with the remote username on the server. For example, if the image file resides in the home directory of a user on the server, specify that user's name as the remote username.

Before you begin downloading or uploading an image file by using RCP, do these tasks:

- Ensure that the workstation acting as the RCP server supports the remote shell (rsh).
- Ensure that the switch has a route to the RCP server. The switch and the server must be in the same subnetwork if you do not have a router to route traffic between subnets. Check connectivity to the RCP server by using the **ping** command.
- If you are accessing the switch through the console or a Telnet session and you do not have a valid username, make sure that the current RCP username is the one that you want to use for the RCP download. You can enter the **show users** privileged EXEC command to view the valid username. If you do not want to use this username, create a new RCP username by using the **ip rcmd remote-username username** global configuration command to be used during all archive operations. The new username is stored in NVRAM. If you are accessing the switch through a Telnet session and you have a valid username, this username is used, and there is no need to set the RCP username. Include the username in the **archive download-sw** or **archive upload-sw** privileged EXEC command if you want to specify a username only for that operation.
- When you upload an image to the RCP to the server, it must be properly configured to accept the RCP write request from the user on the switch. For UNIX systems, you must add an entry to the `.rhosts` file for the remote user on the RCP server. For example, suppose the switch contains these configuration lines:

```
hostname Switch1
ip rcmd remote-username User0
```

If the switch IP address translates to *Switch1.company.com*, the `.rhosts` file for User0 on the RCP server should contain this line:

```
Switch1.company.com Switch1
```

For more information, refer to the documentation for your RCP server.

Downloading an Image File By Using RCP

You can download a new image file and replace or keep the current image.

Beginning in privileged EXEC mode, follow Steps 1 through 6 to download a new image from an RCP server and overwrite the existing image. To keep the current image, skip Step 6.

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using RCP” section on page B-29.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	archive download-sw /overwrite /reload rtp:[[//[username@]location]/directory]/image-name.tar]	Download the image file from the RCP server to the switch, and overwrite the current image. <ul style="list-style-type: none"> • The /overwrite option overwrites the software image in Flash memory with the downloaded image. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For <i>//username</i>, specify the username. For the RCP copy request to execute successfully, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page B-29. • For <i>@location</i>, specify the IP address of the RCP server. • For <i>/directory/image-name.tar</i>, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

Command	Purpose
Step 7 archive download-sw /leave-old-sw /reload rcp:[[/[username@]location]/directory]/image-name.tar]	<p>Download the image file from the RCP server to the switch, and keep the current image.</p> <ul style="list-style-type: none"> • The /leave-old-sw option keeps the old software version after a download. • The /reload option reloads the system after downloading the image unless the configuration has been changed and not been saved. • For //username, specify the username. For the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page B-29. • For @location, specify the IP address of the RCP server. • For /directory]/image-name.tar, specify the directory (optional) and the image to download. Directory and image names are case sensitive.

The download algorithm verifies that the image is appropriate for the switch model and that enough DRAM is present, or it aborts the process and reports an error. If you specify the **/overwrite** option, the download algorithm removes the existing image on the Flash device whether or not it is the same as the new one, downloads the new image, and then reloads the software.

**Note**

If the Flash device has sufficient space to hold two images and you want to overwrite one of these images with the same version, you must specify the **/overwrite** option.

If you specify the **/leave-old-sw**, the existing files are not removed. If there is not enough room to install the new image and keep the running image, the download process stops, and an error message is displayed.

The algorithm installs the downloaded image onto the system board Flash device (flash:). The image is placed into a new directory named with the software version string, and the BOOT environment variable is updated to point to the newly installed image.

If you kept the old software during the download process (you specified the **/leave-old-sw** keyword), you can remove it by entering the **delete /force /recursive filesystem:/file-url** privileged EXEC command. For *filesystem*, use **flash:** for the system board Flash device. For *file-url*, enter the directory name of the old software image. All the files in the directory and the directory are removed.

**Caution**

For the download and upload algorithms to operate properly, do *not* rename image names.

Uploading an Image File By Using RCP

You can upload an image from the switch to an RCP server. You can later download this image to the same switch or to another switch of the same type.

The upload feature should be used only if the HTML pages associated with the Cluster Management Suite (CMS) have been installed with the existing image.

Beginning in privileged EXEC mode, follow these steps to upload an image to an RCP server:

	Command	Purpose
Step 1		Verify that the RCP server is properly configured by referring to the “Preparing to Download or Upload an Image File By Using RCP” section on page B-29.
Step 2		Log into the switch through the console port or a Telnet session.
Step 3	configure terminal	Enter global configuration mode. This step is required only if you override the default remote username (see Steps 4 and 5).
Step 4	ip rcmd remote-username <i>username</i>	(Optional) Specify the remote username.
Step 5	end	Return to privileged EXEC mode.
Step 6	archive upload-sw rmp: [[[//[username@]location]/directory]/image-name.tar]	Upload the currently running switch image to the RCP server. <ul style="list-style-type: none"> For <i>//username</i>, specify the username; for the RCP copy request to execute, an account must be defined on the network server for the remote username. For more information, see the “Preparing to Download or Upload an Image File By Using RCP” section on page B-29. For <i>@location</i>, specify the IP address of the RCP server. For <i>/directory]/image-name.tar</i>, specify the directory (optional) and the name of the software image to be uploaded. Directory and image names are case sensitive. The <i>image-name.tar</i> is the name of software image to be stored on the server.

The **archive upload-sw** privileged EXEC command builds an image file on the server by uploading these files in order: info, the Cisco IOS image, and the HTML files. After these files are uploaded, the upload algorithm creates the tar file format.



Caution

For the download and upload algorithms to operate properly, do *not* rename image names.



Unsupported Commands in Cisco IOS Release 12.1(19)EA1

This appendix lists some of the command-line interface (CLI) commands that appear when you enter the question mark (?) at the Catalyst 2970 switch prompt but are not supported in this release, either because they are not tested, or because of Catalyst 2970 hardware limitations. This is not a complete list. The unsupported commands are listed by software feature and command mode.

Access Control Lists

Unsupported Privileged EXEC Commands

access-enable [host] [timeout *minutes*]

access-template [access-list-number | name] [dynamic-name] [source] [destination] [timeout *minutes*]

clear access-template [access-list-number | name] [dynamic-name] [source] [destination].

Unsupported Global Configuration Commands

access-list rate-limit *acl-index* {*precedence* | **mask** *prec-mask*}

access-list dynamic extended

Unsupported Debug Commands

debug platform cli-redirection main

debug platform configuration

IGMP Snooping Commands

Unsupported Global Configuration Commands

ip igmp snooping source-only-learning
ip igmp snooping tcn

Interface Commands

Unsupported Privileged EXEC Commands

show interfaces [*interface-id* | **vlan** *vlan-id*] [**crb** | **fair-queue** | **irb** | **mac-accounting** | **precedence** | **irb** | **random-detect** | **rate-limit** | **shape**]

Unsupported Global Configuration Commands

interface tunnel

Unsupported Interface Configuration Commands

switchport broadcast *level*
switchport multicast *level*
switchport unicast *level*



Note

These commands have been replaced by the **storm-control** {**broadcast** | **multicast** | **unicast**} **level** *level* [*level*] interface configuration command.

Network Address Translation (NAT) Commands

Unsupported User EXEC Commands

clear ip nat translation
show ip nat statistics
show ip nat translations

Unsupported Global Configuration Commands

`ip nat inside destination`
`ip nat inside source`
`ip nat outside source`
`ip nat pool`

Unsupported Interface Configuration Commands

`ip nat`

RADIUS

Unsupported Global Configuration Commands

`aaa nas port extended`
`radius-server attribute nas-port`
`radius-server configure`
`radius-server extended-portnames`

SNMP

Unsupported Global Configuration Commands

`snmp-server enable informs`
`snmp-server enable traps flash insertion`
`snmp-server enable traps flash removal`
`snmp-server ifindex persist`

Spanning Tree

Unsupported Global Configuration Commands

`spanning-tree etherchannel guard misconfig`
`spanning-tree pathcost method {long | short}`

Unsupported Interface Configuration Commands

spanning-tree stack-port

VLAN

Unsupported vlan-config Commands

private-vlan

Unsupported User EXEC Commands

show running-config vlan

show vlan ifindex

show vlan private-vlan

VTP

Unsupported Privileged EXEC Commands

vtp {password *password* | pruning | version *number*}private-vlan



Note

This command has been replaced by the **vtp** global configuration command.

Miscellaneous

Unsupported Global Configuration Commands

errdisable detect cause dhcp-rate-limit

errdisable recovery cause dhcp-rate-limit

errdisable recovery cause unicast flood

service compress-config



Numerics

802.1D

See STP

802.1Q

and trunk ports [9-3](#)

configuration limitations [11-17](#)

encapsulation [11-15](#)

native VLAN for untagged traffic [11-21](#)

802.1S

See MSTP

802.1W

See RSTP

802.1X

See port-based authentication

802.3AD

See EtherChannel

802.3Z flow control [9-12](#)

A

abbreviating commands [2-4](#)

AC (command switch) [5-10, 5-18](#)

access-class command [26-16](#)

access control entries

See ACEs

access-denied response, VMPS [11-26](#)

access groups

applying ACLs to interfaces [26-17](#)

IP [26-17](#)

accessing

clusters, switch [5-13](#)

command switches [5-11](#)

member switches [5-13](#)

switch clusters [5-13](#)

access lists

See ACLs

access ports

defined [9-2](#)

in switch clusters [5-9](#)

accounting

with RADIUS [7-28](#)

with TACACS+ [7-11, 7-17](#)

ACEs

and QoS [27-7](#)

defined [26-2](#)

Ethernet [26-2](#)

IP [26-2](#)

ACLs

ACEs [26-2](#)

any keyword [26-9](#)

applying

time ranges to [26-13](#)

to an interface [26-16](#)

to QoS [27-7](#)

classifying traffic for QoS [27-37](#)

comments in [26-15](#)

compiling [26-17](#)

defined [26-1, 26-6](#)

examples of [26-17, 27-37](#)

extended IP

configuring for QoS classification [27-38](#)

creating [26-8](#)

matching criteria [26-6](#)

hardware and software handling [26-17](#)

host keyword [26-10](#)

- IP
 - applying to interface [26-16](#)
 - applying to interfaces [26-16](#)
 - creating [26-6](#)
 - fragments and QoS guidelines [27-29](#)
 - implicit deny [26-8, 26-11, 26-13](#)
 - implicit masks [26-8](#)
 - matching criteria [26-6](#)
 - named [26-11](#)
 - terminal lines, setting on [26-16](#)
 - undefined [26-17](#)
- MAC extended [26-20, 27-39](#)
- matching [26-6, 26-17](#)
- monitoring [26-29](#)
- named [26-11](#)
- number per QoS class map [27-29](#)
- numbers [26-6](#)
- port [26-2](#)
- precedence of [26-2](#)
- QoS [27-7, 27-37](#)
- standard IP
 - configuring for QoS classification [27-37](#)
 - creating [26-7](#)
 - matching criteria [26-6](#)
- supported features [26-17](#)
- support for [1-6](#)
- time ranges [26-13](#)
- unsupported features [26-5](#)
- VLAN maps
 - configuration guidelines [26-23](#)
 - configuring [26-22](#)
- address aliasing [18-2](#)
- addresses
 - displaying the MAC address table [6-28](#)
 - dynamic
 - accelerated aging [14-8](#)
 - changing the aging time [6-23](#)
 - default aging [14-8](#)
 - defined [6-22](#)
 - learning [6-22](#)
 - removing [6-24](#)
 - MAC, discovering [6-29](#)
 - multicast, STP address management [14-8](#)
 - static
 - adding and removing [6-26](#)
 - defined [6-22](#)
 - address resolution [6-29](#)
- Address Resolution Protocol
 - See ARP
 - See ARP table
- advertisements
 - CDP [20-1](#)
 - VTP [11-18, 12-3](#)
- aggregated ports
 - See EtherChannel
- aggregate policers [27-45](#)
- aggregate policing [1-7](#)
- aging, accelerating [14-8](#)
- aging time
 - accelerated
 - for MSTP [15-20](#)
 - for STP [14-8, 14-21](#)
 - MAC address table [6-23](#)
 - maximum
 - for MSTP [15-21](#)
 - for STP [14-21](#)
- alarms, RMON [23-3](#)
- allowed-VLAN list [11-20](#)
- ARP
 - defined [1-4, 6-29](#)
- ARP table
 - address resolution [6-29](#)
 - managing [6-29](#)
- attributes, RADIUS
 - vendor-proprietary [7-31](#)
 - vendor-specific [7-29](#)
- audience [xxvii](#)

authentication

- local mode with AAA [7-36](#)
- NTP associations [6-5](#)
- RADIUS
 - key [7-21](#)
 - login [7-23](#)
- See also port-based authentication
- TACACS+
 - defined [7-11](#)
 - key [7-13](#)
 - login [7-14](#)

authoritative time source, described [6-2](#)

authorization

- with RADIUS [7-27](#)
- with TACACS+ [7-11, 7-16](#)

authorized ports with 802.1X [8-4](#)

autoconfiguration [4-3](#)

automatic discovery

- adding member switches [5-16](#)
- considerations
 - beyond a non-candidate device [5-8](#)
 - brand new switches [5-9](#)
 - connectivity [5-5](#)
 - different VLANs [5-7](#)
 - management VLANs [5-8](#)
 - non-CDP-capable devices [5-6](#)
 - non-cluster-capable devices [5-6](#)
- creating a cluster standby group [5-18](#)
- in switch clusters [5-5](#)
- See also CDP

automatic QoS

- See QoS

automatic recovery, clusters [5-10](#)

- See also HSRP

autonegotiation

- duplex mode [1-3](#)
- interface configuration guidelines [9-10](#)
- mismatches [29-12](#)

autosensing, port speed [1-3](#)

auxiliary VLAN

- See voice VLAN

availability, features [1-4](#)

B

BackboneFast

- described [16-5](#)
- enabling [16-13](#)
- support for [1-5](#)

banners

- configuring
 - login [6-21](#)
 - message-of-the-day login [6-20](#)
- default configuration [6-19](#)
- when displayed [6-19](#)

blocking packets [19-6](#)

booting

- boot loader, function of [4-2](#)
- boot process [4-1](#)
- manually [4-13](#)
- specific image [4-14](#)

boot loader

- accessing [4-15](#)
- described [4-2](#)
- environment variables [4-15](#)
- prompt [4-15](#)
- trap-door mechanism [4-2](#)

BPDU

- error-disabled state [16-3](#)
- filtering [16-3](#)
- RSTP format [15-9](#)

BPDU filtering

- described [16-3](#)
- enabling [16-12](#)
- support for [1-5](#)

BPDU guard

- described [16-3](#)
- enabling [16-11](#)
- support for [1-5](#)

bridge protocol data unit

See BPDU

broadcast storm-control command [19-3](#)broadcast storms [19-2](#)

C
cables, monitoring for unidirectional links [21-1](#)

candidate switch

- adding [5-16](#)
- automatic discovery [5-5](#)
- defined [5-4](#)
- HC [5-18](#)
- passwords [5-17](#)
- requirements [5-4](#)
- standby group [5-18](#)

See also command switch, cluster standby group, and member switch

caution, described [xxviii](#)CC (command switch) [5-19](#)

CDP

- and trusted boundary [27-34](#)
- automatic discovery in switch clusters [5-5](#)
- configuring [20-2](#)
- default configuration [20-2](#)
- described [20-1](#)
- disabling for routing device [20-3, 20-4](#)
- enabling and disabling
 - on an interface [20-4](#)
 - on a switch [20-3](#)
- monitoring [20-5](#)
- overview [20-1](#)
- support for [1-4](#)
- transmission timer and holdtime, setting [20-2](#)
- updates [20-2](#)

CGMP

- as IGMP snooping learning method [18-8](#)
- joining multicast group [18-3](#)

Cisco 7960 IP Phone [13-1](#)

Cisco Discovery Protocol

See CDP

Cisco IOS File System

See IFS

CiscoWorks 2000 [1-3, 25-5](#)

class maps for QoS

- configuring [27-40](#)
- described [27-7](#)
- displaying [27-65](#)

class of service

See CoS

clearing interfaces [9-17](#)

CLI

- abbreviating commands [2-4](#)
- command modes [2-1](#)
- described [1-3](#)
- editing features
 - enabling and disabling [2-7](#)
 - keystroke editing [2-7](#)
 - wrapped lines [2-8](#)
- error messages [2-5](#)
- filtering command output [2-9](#)
- getting help [2-3](#)
- history
 - changing the buffer size [2-5](#)
 - described [2-5](#)
 - disabling [2-6](#)
 - recalling commands [2-6](#)
- managing clusters [5-21](#)
- no and default forms of commands [2-4](#)

client mode, VTP [12-3](#)

clock

See system clock

Cluster Management Suite

See CMS

- clusters, switch
 - accessing [5-13](#)
 - adding member switches [5-16](#)
 - automatic discovery [5-5](#)
 - automatic recovery [5-10](#)
 - benefits [1-2](#)
 - command switch configuration [5-16](#)
 - compatibility [5-4](#)
 - creating [5-15](#)
 - creating a cluster standby group [5-18](#)
 - described [5-1](#)
 - LRE profile considerations [5-15](#)
 - managing
 - through CLI [5-21](#)
 - through SNMP [5-21](#)
 - planning [5-4](#)
 - planning considerations
 - automatic discovery [5-5](#)
 - automatic recovery [5-10](#)
 - CLI [5-21](#)
 - host names [5-13](#)
 - IP addresses [5-13](#)
 - LRE profiles [5-15](#)
 - passwords [5-14](#)
 - RADIUS [5-14](#)
 - SNMP [5-14, 5-21](#)
 - switch-specific features [5-15](#)
 - TACACS+ [5-14](#)
 - redundancy [5-18](#)
 - troubleshooting [5-20](#)
 - verifying [5-20](#)
 - See also candidate switch, command switch, cluster standby group, member switch, and standby command switch
- cluster standby group
 - automatic recovery [5-12](#)
 - considerations [5-11](#)
 - creating [5-18](#)
 - defined [5-2](#)
 - requirements [5-3](#)
 - virtual IP address [5-11](#)
 - See also HSRP
- CMS
 - benefits [1-2](#)
 - configuration modes [3-5](#)
 - described [1-2, 1-3](#)
 - Front Panel view
 - described [3-2](#)
 - operating systems and supported browsers [3-8](#)
 - requirements [3-7 to 3-9](#)
 - Topology view [3-14](#)
 - wizards [3-6](#)
 - Coarse Wave Division Multiplexer
 - See CWDM
 - command-line interface
 - See CLI
 - command modes [2-1](#)
 - commands
 - abbreviating [2-4](#)
 - no and default [2-4](#)
 - setting privilege levels [7-8](#)
 - command switch
 - accessing [5-11](#)
 - active (AC) [5-10, 5-18](#)
 - command switch with HSRP disabled (CC) [5-19](#)
 - configuration conflicts [29-11](#)
 - defined [5-2](#)
 - enabling [5-16](#)
 - passive (PC) [5-10, 5-18](#)
 - password privilege levels [5-21](#)
 - priority [5-10](#)
 - recovery
 - from command-switch failure [5-10](#)
 - from failure [29-8](#)
 - from lost member connectivity [29-11](#)
 - redundant [5-10, 5-18](#)

- replacing
 - with another switch [29-10](#)
 - with cluster member [29-8](#)
- requirements [5-3](#)
- standby (SC) [5-10, 5-18](#)
- See also candidate switch, cluster standby group, member switch, and standby command switch
- community strings
 - configuring [5-14, 25-8](#)
 - for cluster switches [25-4](#)
 - in clusters [5-14](#)
 - overview [25-4](#)
 - SNMP [5-14](#)
- config.text [4-12](#)
- configuration conflicts, recovering from lost member connectivity [29-11](#)
- configuration examples, network [1-10](#)
- configuration files
 - clearing the startup configuration [B-19](#)
 - creating using a text editor [B-10](#)
 - default name [4-12](#)
 - deleting a stored configuration [B-19](#)
 - described [B-8](#)
- downloading
 - automatically [4-13](#)
 - preparing [B-10, B-13, B-16](#)
 - reasons for [B-8](#)
 - using FTP [B-13](#)
 - using RCP [B-17](#)
 - using TFTP [B-11](#)
- guidelines for creating and using [B-9](#)
- invalid combinations when copying [B-5](#)
- limiting TFTP server access [25-15](#)
- obtaining with DHCP [4-7](#)
- password recovery disable considerations [7-5](#)
- specifying the filename [4-13](#)
- system contact and location information [25-14](#)
- types and location [B-9](#)
- uploading
 - preparing [B-10, B-13, B-16](#)
 - reasons for [B-8](#)
 - using FTP [B-15](#)
 - using RCP [B-18](#)
 - using TFTP [B-11](#)
- configuration modes, CMS [3-5](#)
- configuration settings, saving [4-11](#)
- configure terminal command [9-5](#)
- config-vlan mode [2-2, 11-6](#)
- conflicts, configuration [29-11](#)
- connections, secure remote [7-38](#)
- connectivity problems [29-13, 29-14, 29-16](#)
- consistency checks in VTP version 2 [12-4](#)
- console port, connecting to [2-10](#)
- conventions
 - command [xxviii](#)
 - for examples [xxviii](#)
 - publication [xxviii](#)
 - text [xxviii](#)
- corrupted software, recovery steps with XMODEM [29-2](#)
- CoS
 - in Layer 2 frames [27-2](#)
 - override priority [13-5](#)
 - trust priority [13-5](#)
- CoS input queue threshold map for QoS [27-14](#)
- CoS output queue threshold map for QoS [27-17](#)
- CoS-to-DSCP map for QoS [27-47](#)
- counters, clearing interface [9-17](#)
- crashinfo file [29-21](#)
- cryptographic software image
 - Kerberos [7-32](#)
 - SSH [7-37](#)
- CWDM [1-15](#)
- CWDM SFPs [1-15](#)

D

daylight saving time [6-14](#)

debugging

enabling all system diagnostics [29-18](#)

enabling for a specific feature [29-18](#)

redirecting error message output [29-19](#)

using commands [29-17](#)

default commands [2-4](#)

default configuration

802.1X [8-10](#)

auto-QoS [27-18](#)

banners [6-19](#)

booting [4-12](#)

CDP [20-2](#)

DHCP [17-3](#)

DNS [6-18](#)

EtherChannel [28-9](#)

IGMP filtering [18-21](#)

IGMP snooping [18-7](#)

IGMP throttling [18-21](#)

initial switch information [4-3](#)

Layer 2 interfaces [9-9](#)

MAC address table [6-23](#)

MSTP [15-12](#)

MVR [18-16](#)

NTP [6-4](#)

optional spanning-tree features [16-9](#)

password and privilege level [7-2](#)

RADIUS [7-20](#)

RMON [23-3](#)

RSPAN [22-9](#)

SNMP [25-7](#)

SPAN [22-9](#)

standard QoS [27-27](#)

STP [14-11](#)

system message logging [24-3](#)

system name and prompt [6-16](#)

TACACS+ [7-13](#)

UDLD [21-4](#)

VLAN, Layer 2 Ethernet interfaces [11-17](#)

VLANs [11-7](#)

VMPS [11-27](#)

voice VLAN [13-3](#)

VTP [12-6](#)

default gateway [4-10](#)

deleting VLANs [11-10](#)

description command [9-14](#)

designing your network, examples [1-10](#)

destination addresses, in ACLs [26-9](#)

destination-IP address based forwarding,
EtherChannel [28-7](#)

destination-MAC address forwarding, EtherChannel [28-7](#)

detecting indirect link failures, STP [16-6](#)

device discovery protocol [20-1](#)

Device Manager [3-14](#)

See also Switch Manager

DHCP-based autoconfiguration

client request message exchange [4-4](#)

configuring

client side [4-3](#)

DNS [4-6](#)

relay device [4-6](#)

server-side [4-5](#)

TFTP server [4-6](#)

example [4-8](#)

lease options

for IP address information [4-5](#)

for receiving the configuration file [4-5](#)

overview [4-3](#)

relationship to BOOTP [4-4](#)

relay support [1-4](#)

support for [1-4](#)

DHCP option 82

configuration guidelines [17-3](#)

default configuration [17-3](#)

displaying [17-5](#)

overview [17-2](#)

DHCP snooping

- configuration guidelines [17-3](#)
- default configuration [17-3](#)
- displaying binding tables [17-5](#)
- displaying configuration [17-5](#)
- message exchange process [17-2](#)
- option 82 data insertion [17-2](#)

Differentiated Services architecture, QoS [27-1](#)Differentiated Services Code Point [27-2](#)directed unicast requests [1-4](#)

directories

- changing [B-3](#)
- creating and removing [B-4](#)
- displaying the working [B-3](#)

discovery, clusters

- See automatic discovery

DNS

- and DHCP-based autoconfiguration [4-6](#)
- default configuration [6-18](#)
- displaying the configuration [6-19](#)
- overview [6-17](#)
- setting up [6-18](#)
- support for [1-4](#)

documentation

- related [xxix](#)

document conventions [xxviii](#)

domain names

- DNS [6-17](#)
- VTP [12-7](#)

Domain Name System

- See DNS

downloading

- configuration files
 - preparing [B-10, B-13, B-16](#)
 - reasons for [B-8](#)
 - using FTP [B-13](#)
 - using RCP [B-17](#)
 - using TFTP [B-11](#)

image files

- deleting old image [B-24](#)
- preparing [B-22, B-25, B-29](#)
- reasons for [B-20](#)
- using FTP [B-26](#)
- using RCP [B-31](#)
- using TFTP [B-23](#)

DSCP [1-6, 27-2](#)DSCP input queue threshold map for QoS [27-14](#)DSCP output queue threshold map for QoS [27-17](#)DSCP-to-CoS map for QoS [27-50](#)DSCP-to-DSCP-mutation map for QoS [27-51](#)DTP [1-5, 11-16](#)

dynamic access ports

- characteristics [11-4](#)
- configuring [11-29](#)
- defined [9-3](#)

dynamic addresses

- See addresses

dynamic auto trunking mode [11-16](#)dynamic desirable trunking mode [11-16](#)

Dynamic Host Configuration Protocol

- See DHCP-based autoconfiguration

dynamic port VLAN membership

- described [11-27](#)
- reconfirming [11-29](#)
- troubleshooting [11-31](#)
- types of connections [11-29](#)

Dynamic Trunking Protocol

- See DTP

E

editing features

- enabling and disabling [2-7](#)
- keystrokes used [2-7](#)
- wrapped lines [2-8](#)

enable password [7-4](#)enable secret password [7-4](#)

- encryption for passwords 7-4
 - environment variables
 - function of 4-16
 - error messages during command entry 2-5
 - EtherChannel
 - 802.3AD, described 28-5
 - automatic creation of 28-3, 28-5
 - channel groups
 - binding physical and logical interfaces 28-3
 - numbering of 28-3
 - configuration guidelines 28-9
 - configuring Layer 2 interfaces 28-10
 - default configuration 28-9
 - described 28-2
 - displaying status 28-17
 - forwarding methods 28-6, 28-12
 - interaction
 - with STP 28-9
 - with VLANs 28-10
 - LACP
 - described 28-5
 - displaying status 28-17
 - hot-standby ports 28-15
 - interaction with other features 28-6
 - modes 28-6
 - port priority 28-16
 - system priority 28-15
 - LACP, support for 1-3
 - load balancing 28-6, 28-12
 - number of interfaces per 28-2
 - PAgP
 - aggregate-port learners 28-13
 - compatibility with Catalyst 1900 28-13
 - described 28-3
 - displaying status 28-17
 - interaction with other features 28-5
 - learn method and priority configuration 28-13
 - modes 28-4
 - silent mode 28-5
 - support for 1-3
 - port-channel interfaces
 - described 28-3
 - numbering of 28-3
 - port groups 9-4
 - support for 1-3
 - Ethernet VLANs
 - adding 11-8
 - defaults and ranges 11-8
 - modifying 11-8
 - events, RMON 23-3
 - examples
 - conventions for xxviii
 - network configuration 1-10
 - expedite queue for QoS
 - configuring 27-63
 - expert mode 3-6
 - Express Setup 3-11
 - extended-range VLANs
 - configuration guidelines 11-12
 - configuring 11-12
 - creating 11-12, 11-13
 - defined 11-1
 - extended system ID
 - MSTP 15-14
 - STP 14-4, 14-14
 - Extensible Authentication Protocol over LAN 8-1
-
- ## F
- fiber-optic, detecting unidirectional links 21-1
 - files
 - copying B-4
 - crashinfo
 - description 29-21

- displaying the contents of [29-21](#)
- location [29-21](#)
- deleting [B-5](#)
- displaying the contents of [B-7](#)
- tar
 - creating [B-6](#)
 - displaying the contents of [B-6](#)
 - extracting [B-7](#)
 - image file format [B-21](#)
- file system
 - displaying available file systems [B-2](#)
 - displaying file information [B-3](#)
 - local file system names [B-1](#)
 - network file system names [B-4](#)
 - setting the default [B-3](#)
- filtering
 - in a VLAN [26-22](#)
 - non-IP traffic [26-20](#)
 - show and more command output [2-9](#)
- filtering show and more command output [2-9](#)
- filters, IP
 - See ACLs, IP
- Flash device, number of [B-1](#)
- flooded traffic, blocking [19-6](#)
- flow-based packet classification [1-6](#)
- flowcharts
 - QoS classification [27-6](#)
 - QoS egress queueing and scheduling [27-15](#)
 - QoS ingress queueing and scheduling [27-13](#)
 - QoS policing and marking [27-9](#)
- flow control [1-3, 9-12](#)
- forward-delay time
 - MSTP [15-20](#)
 - STP [14-21](#)
- FTP
 - accessing MIB files [A-3](#)
 - configuration files
 - downloading [B-13](#)
 - overview [B-12](#)

- preparing the server [B-13](#)
- uploading [B-15](#)
- image files
 - deleting old image [B-28](#)
 - downloading [B-26](#)
 - preparing the server [B-25](#)
 - uploading [B-28](#)

G

- get-bulk-request operation [25-3](#)
- get-next-request operation [25-3, 25-5](#)
- get-request operation [25-3, 25-5](#)
- get-response operation [25-3](#)
- global configuration mode [2-2](#)
- guest VLAN and 802.1X [8-8](#)
- guide
 - audience [xxvii](#)
 - purpose of [xxvii](#)
- guide mode [1-2, 3-5](#)

H

- HC (candidate switch) [5-18](#)
- hello time
 - MSTP [15-19](#)
 - STP [14-20](#)
- help, for the command line [2-3](#)
- history
 - changing the buffer size [2-5](#)
 - described [2-5](#)
 - disabling [2-6](#)
 - recalling commands [2-6](#)
- history table, level and number of syslog messages [24-9](#)
- host names
 - abbreviations appended to [5-18](#)
 - in clusters [5-13](#)
- hosts, limit on dynamic ports [11-31](#)

HP OpenView 1-3

HSRP

automatic cluster recovery 5-12

cluster standby group considerations 5-11

See also clusters, cluster standby group, and standby command switch

ICMP

time exceeded messages 29-16

traceroute and 29-16

ICMP ping

executing 29-13

overview 29-13

IDS appliances

and ingress RSPAN 22-20

and ingress SPAN 22-13

IEEE 802.1P 13-1

ifIndex values, SNMP 25-6

IFS 1-4

IGMP

joining multicast group 18-3

join messages 18-3

leave processing, enabling 18-10

leaving multicast group 18-5

queries 18-4

report suppression

described 18-6

disabling 18-11

support for 1-3

IGMP filtering

configuring 18-21

default configuration 18-21

described 18-20

monitoring 18-25

support for 1-3

IGMP groups

configuring the filtering action 18-24

setting the maximum number 18-23

IGMP profile

applying 18-22

configuration mode 18-21

configuring 18-21

IGMP snooping

and address aliasing 18-2

configuring 18-6

default configuration 18-7

definition 18-2

enabling and disabling 18-7

global configuration 18-7

Immediate Leave 18-6

method 18-8

monitoring 18-12

support for 1-3

VLAN configuration 18-7

IGMP throttling

configuring 18-24

default configuration 18-21

described 18-20

displaying action 18-25

Immediate-Leave, IGMP 18-6

initial configuration

defaults 1-8

See also hardware installation guide

interface

number 9-5

range macros 9-7

interface command 9-5

interface configuration mode 2-3

interfaces

configuration guidelines 9-10

configuring 9-5

configuring speed 9-10

counters, clearing 9-17

described 9-14

- descriptive name, adding [9-14](#)
 - displaying information about [9-16](#)
 - flow control [9-12](#)
 - management [1-3](#)
 - monitoring [9-16](#)
 - naming [9-14](#)
 - physical, identifying [9-5](#)
 - range of [9-6](#)
 - restarting [9-17](#)
 - shutting down [9-17](#)
 - supported [9-5](#)
 - types of [9-1](#)
 - interfaces range macro command [9-7](#)
 - interface types [9-5](#)
 - Inter-Switch Link
 - See ISL
 - Intrusion Detection System
 - See IDS
 - inventory, cluster [5-20](#)
 - IOS File System
 - See IFS
 - ip access group command [26-17](#)
 - IP ACLs
 - applying to an interface [26-16](#)
 - extended, creating [26-8](#)
 - for QoS classification [27-7](#)
 - implicit deny [26-8, 26-11, 26-13](#)
 - implicit masks [26-8](#)
 - named [26-11](#)
 - standard, creating [26-7](#)
 - undefined [26-17](#)
 - virtual terminal lines, setting on [26-16](#)
 - IP addresses
 - candidate or member [5-4, 5-13](#)
 - cluster access [5-2](#)
 - command switch [5-3, 5-11, 5-13](#)
 - discovering [6-29](#)
 - redundant clusters [5-11](#)
 - standby command switch [5-11, 5-13](#)
 - See also IP information
 - ip igmp profile command [18-21](#)
 - IP information
 - assigned
 - manually [4-10](#)
 - through DHCP-based autoconfiguration [4-3](#)
 - default configuration [4-3](#)
 - IP phones
 - and QoS [13-1](#)
 - automatic classification and queueing [27-18](#)
 - configuring [13-4](#)
 - ensuring port security with QoS [27-34](#)
 - trusted boundary for QoS [27-34](#)
 - IP precedence [27-2](#)
 - IP-precedence-to-DSCP map for QoS [27-48](#)
 - IP protocols in ACLs [26-9](#)
 - IP traceroute
 - executing [29-16](#)
 - overview [29-16](#)
 - ISL
 - and trunk ports [9-3](#)
 - encapsulation [1-5, 11-15](#)
-
- ## J
- join messages, IGMP [18-3](#)
-
- ## K
- KDC
 - described [7-32](#)
 - See also Kerberos

Kerberos

- authenticating to
 - boundary switch [7-35](#)
 - KDC [7-35](#)
 - network services [7-35](#)
 - configuration examples [7-32](#)
 - configuring [7-36](#)
 - credentials [7-32](#)
 - cryptographic software image [7-32](#)
 - described [7-32](#)
 - KDC [7-32](#)
 - operation [7-34](#)
 - realm [7-33](#)
 - server [7-33](#)
 - support for [1-6](#)
 - switch as trusted third party [7-32](#)
 - terms [7-33](#)
 - TGT [7-34](#)
 - tickets [7-32](#)
- key distribution center
- See KDC

L

LACP

- See EtherChannel

Layer 2 frames, classification with CoS [27-2](#)

Layer 2 interfaces, default configuration [9-9](#)

Layer 2 traceroute

- and ARP [29-15](#)
- and CDP [29-14](#)
- described [29-14](#)
- IP addresses and subnets [29-15](#)
- MAC addresses and VLANs [29-15](#)
- multicast traffic [29-15](#)
- multiple devices on a port [29-15](#)
- unicast traffic [29-14](#)
- usage guidelines [29-14](#)

Layer 2 trunks [11-15](#)

Layer 3 packets, classification methods [27-2](#)

leave processing, IGMP [18-10](#)

LEDs, switch

- See hardware installation guide

line configuration mode [2-3](#)

Link Aggregation Control Protocol

- See EtherChannel
- See LACP

links, unidirectional [21-1](#)

login authentication

- with RADIUS [7-23](#)
- with TACACS+ [7-14](#)

login banners [6-19](#)

log messages

- See system message logging

Long-Reach Ethernet (LRE) technology [1-11](#)

loop guard

- described [16-8](#)
- enabling [16-15](#)
- support for [1-5](#)

LRE profiles, considerations in switch clusters [5-15](#)

M

MAC addresses

- aging time [6-23](#)
- and VLAN association [6-23](#)
- building the address table [6-22](#)
- default configuration [6-23](#)
- discovering [6-29](#)
- displaying [6-28](#)
- displaying in DHCP snooping binding table [17-5](#)
- dynamic
 - learning [6-22](#)
 - removing [6-24](#)
- in ACLs [26-20](#)
- static
 - adding [6-27](#)
 - allowing [6-28](#)

- characteristics of [6-26](#)
 - dropping [6-28](#)
 - removing [6-27](#)
- MAC address notification, support for [1-7](#)
- MAC address-to-VLAN mapping [11-26](#)
- MAC extended access lists
 - applying to Layer 2 interfaces [26-21](#)
 - configuring for QoS [27-39](#)
 - creating [26-20](#)
 - defined [26-20](#)
 - for QoS classification [27-5](#)
- macros
 - See SmartPort macros
- manageability features [1-4](#)
- management access
 - in-band
 - browser session [1-4](#)
 - CLI session [1-4](#)
 - SNMP [1-4](#)
 - out-of-band console port connection [1-4](#)
- management options
 - benefits
 - clustering [1-2](#)
 - CMS [1-2](#)
 - CLI [2-1](#)
 - overview [1-3](#)
- management VLAN
 - considerations in switch clusters [5-8](#)
 - discovery through different management VLANs [5-8](#)
- mapping tables for QoS
 - configuring
 - CoS-to-DSCP [27-47](#)
 - DSCP [27-47](#)
 - DSCP-to-CoS [27-50](#)
 - DSCP-to-DSCP-mutation [27-51](#)
 - IP-precedence-to-DSCP [27-48](#)
 - policed-DSCP [27-49](#)
 - described [27-10](#)
- marking
 - action in policy map [27-42](#)
 - action with aggregate policers [27-45](#)
 - described [27-3, 27-8](#)
- matching, ACLs [26-6](#)
- maximum aging time
 - MSTP [15-21](#)
 - STP [14-21](#)
- maximum hop count, MSTP [15-21](#)
- membership mode, VLAN port [11-3](#)
- member switch
 - adding [5-16](#)
 - automatic discovery [5-5](#)
 - defined [5-2](#)
 - managing [5-21](#)
 - passwords [5-13](#)
 - recovering from lost connectivity [29-11](#)
 - requirements [5-4](#)
 - See also candidate switch, cluster standby group, and standby command switch
- menu bar
 - variations [3-4](#)
- messages
 - to users through banners [6-19](#)
- MIBs
 - accessing files with FTP [A-3](#)
 - location of files [A-3](#)
 - overview [25-1](#)
 - SNMP interaction with [25-5](#)
 - supported [A-1](#)
- mirroring traffic for analysis [22-1](#)
- mismatches, autonegotiation [29-12](#)
- module number [9-5](#)
- monitoring
 - access groups [26-29](#)
 - ACL configuration [26-29](#)
 - cables for unidirectional links [21-1](#)
 - CDP [20-5](#)
 - features [1-7](#)

- IGMP
 - filters [18-25](#)
 - snooping [18-12](#)
- interfaces [9-16](#)
- multicast router interfaces [18-12](#)
- MVR [18-19](#)
- network traffic for analysis with probe [22-2](#)
- port
 - blocking [19-15](#)
 - protection [19-15](#)
- speed and duplex mode [9-11](#)
- traffic flowing among switches [23-1](#)
- traffic suppression [19-15](#)
- VLAN
 - filters [26-29](#)
 - maps [26-29](#)
- VLANs [11-14](#)
- VMPS [11-30](#)
- VTP [12-15](#)
- MSTP
 - boundary ports
 - configuration guidelines [15-13](#)
 - described [15-5](#)
 - BPDU filtering
 - described [16-3](#)
 - enabling [16-12](#)
 - BPDU guard
 - described [16-3](#)
 - enabling [16-11](#)
 - CIST, described [15-3](#)
 - configuration guidelines [15-12, 16-9](#)
 - configuring
 - forward-delay time [15-20](#)
 - hello time [15-19](#)
 - link type for rapid convergence [15-22](#)
 - maximum aging time [15-21](#)
 - maximum hop count [15-21](#)
 - MST region [15-13](#)
 - path cost [15-18](#)
 - port priority [15-17](#)
 - root switch [15-14](#)
 - secondary root switch [15-16](#)
 - switch priority [15-19](#)
- CST
 - defined [15-3](#)
 - operations between regions [15-4](#)
- default configuration [15-12](#)
- default optional feature configuration [16-9](#)
- displaying status [15-23](#)
- enabling the mode [15-13](#)
- extended system ID
 - effects on root switch [15-14](#)
 - effects on secondary root switch [15-16](#)
 - unexpected behavior [15-15](#)
- instances supported [14-9](#)
- interface state, blocking to forwarding [16-2](#)
- interoperability and compatibility among modes [14-10](#)
- interoperability with 802.1D
 - described [15-5](#)
 - restarting migration process [15-22](#)
- IST
 - defined [15-3](#)
 - master [15-3](#)
 - operations within a region [15-3](#)
- loop guard
 - described [16-8](#)
 - enabling [16-15](#)
- mapping VLANs to MST instance [15-13](#)
- MST region
 - CIST [15-3](#)
 - configuring [15-13](#)
 - described [15-2](#)
 - hop-count mechanism [15-5](#)
 - IST [15-3](#)
 - supported spanning-tree instances [15-2](#)
- optional features supported [1-5](#)
- overview [15-2](#)

Port Fast

- described [16-2](#)
- enabling [16-10](#)

preventing root switch selection [16-7](#)

root guard

- described [16-7](#)
- enabling [16-14](#)

root switch

- configuring [15-15](#)
- effects of extended system ID [15-14](#)
- unexpected behavior [15-15](#)

shutdown Port Fast-enabled port [16-3](#)

status, displaying [15-23](#)

multicast groups

- Immediate Leave [18-6](#)
- joining [18-3](#)
- leaving [18-5](#)
- static joins [18-10](#)

multicast packets, blocking [19-6](#)

multicast router interfaces, monitoring [18-12](#)

multicast router ports, adding [18-9](#)

multicast storm-control command [19-4](#)

multicast storms [19-2](#)

Multicast VLAN Registration

- See MVR

Multiple Spanning Tree Protocol

- See MSTP

MVR

- and address aliasing [18-16](#)
- configuring interfaces [18-18](#)
- default configuration [18-16](#)
- described [18-13](#)
- modes [18-17](#)
- monitoring [18-19](#)
- setting global parameters [18-16](#)
- support for [1-3](#)

N

named IP ACLs [26-11](#)

native VLAN

- configuring [11-21](#)
- default [11-21](#)

network configuration examples

- increasing network performance [1-10](#)
- long-distance, high-bandwidth transport [1-15](#)
- providing network services [1-11](#)
- server aggregation and Linux server cluster [1-13](#)
- small to medium-sized network [1-14](#)

network design

- performance [1-10](#)
- services [1-11](#)

network management

- CDP [20-1](#)
- RMON [23-1](#)
- SNMP [25-1](#)

Network Time Protocol

- See NTP

no commands [2-4](#)

non-IP traffic filtering [26-20](#)

nontrunking mode [11-16](#)

normal-range VLANs

- configuration modes [11-6](#)
- defined [11-1](#)

note, described [xxviii](#)

NTP

- associations
 - authenticating [6-5](#)
 - defined [6-2](#)
 - enabling broadcast messages [6-7](#)
 - peer [6-6](#)
 - server [6-6](#)
- default configuration [6-4](#)
- displaying the configuration [6-11](#)

- overview [6-2](#)
- restricting access
 - creating an access group [6-9](#)
 - disabling NTP services per interface [6-10](#)
- source IP address, configuring [6-10](#)
- stratum [6-2](#)
- support for [1-4](#)
- synchronizing devices [6-6](#)
- time
 - services [6-2](#)
 - synchronizing [6-2](#)

O

- options, management [1-3](#)
- out-of-profile markdown [1-7](#)

P

- packet modification, with QoS [27-17](#)
- PAgP
 - See EtherChannel
- passwords
 - default configuration [7-2](#)
 - disabling recovery of [7-5](#)
 - encrypting [7-4](#)
 - for security [1-5](#)
 - in clusters [5-14, 5-17](#)
 - overview [7-1](#)
 - recovery of [29-4](#)
 - setting
 - enable [7-3](#)
 - enable secret [7-4](#)
 - Telnet [7-6](#)
 - with usernames [7-7](#)
 - VTP domain [12-8](#)

- path cost
 - MSTP [15-18](#)
 - STP [14-18](#)
- PC (passive command switch) [5-10, 5-18](#)
- performance, network design [1-10](#)
- performance features [1-3](#)
- per-VLAN spanning-tree plus
 - See PVST+
- physical ports [9-2](#)
- PIM-DVMRP, as snooping method [18-8](#)
- ping
 - character output description [29-13](#)
 - executing [29-13](#)
 - overview [29-13](#)
- policed-DSCP map for QoS [27-49](#)
- policers
 - configuring
 - for each matched traffic class [27-42](#)
 - for more than one traffic class [27-45](#)
 - described [27-3](#)
 - displaying [27-65](#)
 - number of [27-29](#)
 - types of [27-8](#)
- policing
 - described [27-3](#)
 - token-bucket algorithm [27-9](#)
- policy maps for QoS
 - characteristics of [27-42](#)
 - configuring [27-42](#)
 - described [27-7](#)
 - displaying [27-65](#)
- port ACLs
 - defined [26-2](#)
 - types of [26-2](#)
- Port Aggregation Protocol
 - See EtherChannel
 - See PAgP

- port-based authentication
 - authentication server
 - defined 8-2
 - RADIUS server 8-2
 - client, defined 8-2
 - configuration guidelines 8-11
 - configuring
 - 802.1X authentication 8-12
 - guest VLAN 8-18
 - host mode 8-17
 - manual re-authentication of a client 8-15
 - periodic re-authentication 8-14
 - quiet period 8-15
 - RADIUS server 8-14
 - RADIUS server parameters on the switch 8-13
 - switch-to-client frame-retransmission number 8-17
 - switch-to-client retransmission time 8-16
 - default configuration 8-10
 - described 8-1
 - device roles 8-2
 - displaying statistics 8-19
 - EAPOL-start frame 8-3
 - EAP-request/identity frame 8-3
 - EAP-response/identity frame 8-3
 - encapsulation 8-2
 - guest VLAN
 - configuration guidelines 8-8
 - described 8-8
 - initiation and message exchange 8-3
 - method lists 8-12
 - multiple-hosts mode, described 8-17
 - per-user ACLs
 - AAA authorization 8-12
 - configuration tasks 8-9
 - described 8-9
 - RADIUS server attributes 8-9
 - ports
 - authorized and unauthorized 8-4
 - voice VLAN 8-6
 - port security
 - and voice VLAN 8-6
 - described 8-6
 - interactions 8-6
 - multiple-hosts mode 8-18
 - resetting to default values 8-19
 - statistics, displaying 8-19
 - switch
 - as proxy 8-2
 - RADIUS client 8-2
 - topologies, supported 8-5
 - upgrading from a previous release 8-12
 - VLAN assignment
 - AAA authorization 8-12
 - characteristics 8-7
 - configuration tasks 8-8
 - described 8-7
 - voice VLAN
 - described 8-6
 - PVID 8-6
 - VVID 8-6
- port blocking 1-3, 19-6
- port-channel
 - See EtherChannel
- Port Fast
 - described 16-2
 - enabling 16-10
 - mode, spanning tree 11-27
 - support for 1-5
- port membership modes, VLAN 11-3
- port priority
 - MSTP 15-17
 - STP 14-17
- ports
 - access 9-2
 - blocking 19-6
 - dynamic access 11-4

- protected [19-5](#)
 - secure [19-7](#)
 - static-access [11-3, 11-11](#)
 - switch [9-2](#)
 - trunks [11-3, 11-15](#)
 - VLAN assignments [11-11](#)
 - port security
 - aging [19-13](#)
 - and QoS trusted boundary [27-34](#)
 - configuring [19-10](#)
 - default configuration [19-9](#)
 - described [19-7](#)
 - displaying [19-15](#)
 - on trunk ports [19-11](#)
 - sticky learning [19-8](#)
 - violations [19-8](#)
 - with other features [19-9](#)
 - port-shutdown response, VMPS [11-26](#)
 - preferential treatment of traffic
 - See QoS
 - preventing unauthorized access [7-1](#)
 - priority
 - overriding CoS [13-5](#)
 - trusting CoS [13-5](#)
 - private VLAN edge ports
 - See protected ports
 - privileged EXEC mode [2-2](#)
 - privilege levels
 - changing the default for lines [7-9](#)
 - command switch [5-21](#)
 - exiting [7-10](#)
 - logging into [7-10](#)
 - mapping on member switches [5-21](#)
 - overview [7-2, 7-8](#)
 - setting a command with [7-8](#)
 - protected ports [1-6, 19-5](#)
 - pruning, VTP
 - enabling [12-13](#)
 - enabling on a port [11-21](#)
 - examples [12-5](#)
 - overview [12-4](#)
 - pruning-eligible list
 - changing [11-21](#)
 - for VTP pruning [12-4](#)
 - VLANs [12-14](#)
 - PVST+
 - 802.1Q trunking interoperability [14-10](#)
 - described [14-9](#)
 - instances supported [14-9](#)
-
- ## Q
- QoS
 - auto-QoS
 - categorizing traffic [27-18](#)
 - configuration and defaults display [27-26](#)
 - configuration guidelines [27-22](#)
 - described [27-18](#)
 - disabling [27-23](#)
 - displaying generated commands [27-23](#)
 - displaying the initial configuration [27-26](#)
 - effects on running configuration [27-22](#)
 - egress queue defaults [27-19](#)
 - enabling for VoIP [27-23](#)
 - example configuration [27-24](#)
 - ingress queue defaults [27-19](#)
 - list of generated commands [27-20](#)
 - basic model [27-3](#)
 - classification
 - class maps, described [27-7](#)
 - defined [27-3](#)
 - flowchart [27-6](#)
 - forwarding treatment [27-3](#)
 - in frames and packets [27-2](#)
 - IP ACLs, described [27-5, 27-7](#)
 - MAC ACLs, described [27-5, 27-7](#)
 - options for IP traffic [27-5](#)
 - options for non-IP traffic [27-5](#)

- policy maps, described [27-7](#)
- trust DSCP, described [27-5](#)
- trusted CoS, described [27-5](#)
- trust IP precedence, described [27-5](#)
- class maps
 - configuring [27-40](#)
 - displaying [27-65](#)
- configuration guidelines
 - auto-QoS [27-22](#)
 - standard QoS [27-29](#)
- configuring
 - aggregate policers [27-45](#)
 - auto-QoS [27-18](#)
 - default port CoS value [27-33](#)
 - DSCP maps [27-47](#)
 - DSCP trust states bordering another domain [27-35](#)
 - egress queue characteristics [27-57](#)
 - ingress queue characteristics [27-52](#)
 - IP extended ACLs [27-38](#)
 - IP standard ACLs [27-37](#)
 - MAC ACLs [27-39](#)
 - policy maps [27-42](#)
 - port trust states within the domain [27-31](#)
 - trusted boundary [27-34](#)
- default auto configuration [27-18](#)
- default standard configuration [27-27](#)
- displaying statistics [27-65](#)
- egress queues
 - allocating buffer space [27-58](#)
 - buffer allocation scheme, described [27-16](#)
 - configuring shaped weights for SRR [27-61](#)
 - configuring shared weights for SRR [27-62](#)
 - described [27-4](#)
 - displaying the threshold map [27-60](#)
 - flowchart [27-15](#)
 - mapping DSCP or CoS values [27-60](#)
 - scheduling, described [27-4](#)
 - setting WTD thresholds [27-58](#)
 - WTD, described [27-17](#)
- enabling globally [27-30](#)
- flowcharts
 - classification [27-6](#)
 - egress queueing and scheduling [27-15](#)
 - ingress queueing and scheduling [27-13](#)
 - policing and marking [27-9](#)
- implicit deny [27-7](#)
- ingress queues
 - allocating bandwidth [27-55](#)
 - allocating buffer space [27-54](#)
 - buffer and bandwidth allocation, described [27-14](#)
 - configuring shared weights for SRR [27-55](#)
 - configuring the priority queue [27-56](#)
 - described [27-3](#)
 - displaying the threshold map [27-53](#)
 - flowchart [27-13](#)
 - mapping DSCP or CoS values [27-53](#)
 - priority queue, described [27-14](#)
 - scheduling, described [27-3](#)
 - setting WTD thresholds [27-53](#)
 - WTD, described [27-14](#)
- IP phones
 - automatic classification and queueing [27-18](#)
 - detection and trusted settings [27-18, 27-34](#)
- limiting bandwidth on egress interface [27-64](#)
- mapping tables
 - CoS-to-DSCP [27-47](#)
 - displaying [27-65](#)
 - DSCP-to-CoS [27-50](#)
 - DSCP-to-DSCP-mutation [27-51](#)
 - IP-precedence-to-DSCP [27-48](#)
 - policed-DSCP [27-49](#)
 - types of [27-10](#)

- marked-down actions [27-43](#)
- marking, described [27-3, 27-8](#)
- overview [27-1](#)
- packet modification [27-17](#)
- policers
 - configuring [27-43, 27-45](#)
 - described [27-8](#)
 - displaying [27-65](#)
 - number of [27-29](#)
 - types of [27-8](#)
- policies, attaching to an interface [27-9](#)
- policing
 - described [27-3, 27-8](#)
 - token bucket algorithm [27-9](#)
- policy maps
 - characteristics of [27-42](#)
 - configuring [27-42](#)
 - displaying [27-65](#)
- QoS label, defined [27-3](#)
- queues
 - configuring egress characteristics [27-57](#)
 - configuring ingress characteristics [27-52](#)
 - high priority (expedite) [27-17, 27-63](#)
 - location of [27-11](#)
 - SRR, described [27-12](#)
 - WTD, described [27-11](#)
- rewrites [27-17](#)
- support for [1-6](#)
- trust states
 - bordering another domain [27-35](#)
 - described [27-5](#)
 - trusted device [27-34](#)
 - within the domain [27-31](#)
- quality of service
 - See QoS
- queries, IGMP [18-4](#)

R

RADIUS

- attributes
 - vendor-proprietary [7-31](#)
 - vendor-specific [7-29](#)
- configuring
 - accounting [7-28](#)
 - authentication [7-23](#)
 - authorization [7-27](#)
 - communication, global [7-21, 7-29](#)
 - communication, per-server [7-21](#)
 - multiple UDP ports [7-21](#)
- default configuration [7-20](#)
- defining AAA server groups [7-25](#)
- displaying the configuration [7-31](#)
- identifying the server [7-21](#)
- in clusters [5-14](#)
- limiting the services to the user [7-27](#)
- method list, defined [7-20](#)
- operation of [7-19](#)
- overview [7-18](#)
- suggested network environments [7-18](#)
- support for [1-6](#)
- tracking services accessed by user [7-28](#)
- range
 - macro [9-7](#)
 - of interfaces [9-6](#)
- rapid convergence [15-7](#)
- rapid per-VLAN spanning-tree plus
 - See rapid PVST+
- rapid PVST+
 - 802.1Q trunking interoperability [14-10](#)
 - described [14-9](#)
 - instances supported [14-9](#)
- Rapid Spanning Tree Protocol
 - See RSTP
- rcommand command [5-21](#)

RCP

configuration files

downloading [B-17](#)overview [B-16](#)preparing the server [B-16](#)uploading [B-18](#)

image files

deleting old image [B-32](#)downloading [B-31](#)preparing the server [B-29](#)uploading [B-33](#)reconfirmation interval, VMPS, changing [11-29](#)recovery procedures [29-1](#)

redundancy

EtherChannel [28-2](#)

STP

backbone [14-8](#)path cost [11-24](#)port priority [11-22](#)

redundant clusters

See cluster standby group

redundant links and UplinkFast [16-13](#)reloading software [4-17](#)

Remote Authentication Dial-In User Service

See RADIUS

Remote Copy Protocol

See RCP

Remote Network Monitoring

See RMON

Remote SPAN

See RSPAN

report suppression, IGMP

described [18-6](#)disabling [18-11](#)resetting a UDLD-shutdown interface [21-6](#)

restricting access

NTP services [6-8](#)overview [7-1](#)passwords and privilege levels [7-2](#)RADIUS [7-18](#)TACACS+ [7-10](#)retry count, VMPS, changing [11-30](#)

RFC

1112, IP multicast and IGMP [18-2](#)1157, SNMPv1 [25-2](#)1305, NTP [6-2](#)1757, RMON [23-2](#)1901, SNMPv2C [25-2](#)1902 to 1907, SNMPv2 [25-2](#)2236, IP multicast and IGMP [18-2](#)2273-2275, SNMPv3 [25-2](#)

RMON

default configuration [23-3](#)displaying status [23-6](#)enabling alarms and events [23-3](#)groups supported [23-2](#)overview [23-1](#)

statistics

collecting group Ethernet [23-6](#)collecting group history [23-5](#)support for [1-7](#)

root guard

described [16-7](#)enabling [16-14](#)support for [1-5](#)

root switch

MSTP [15-14](#)STP [14-14](#)

RSPAN

characteristics [22-8](#)configuration guidelines [22-16](#)default configuration [22-9](#)destination ports [22-7](#)displaying status [22-23](#)interaction with other features [22-8](#)monitored ports [22-5](#)monitoring ports [22-7](#)overview [1-7, 22-1](#)

- received traffic [22-4](#)
- session limits [22-10](#)
- sessions
 - creating [22-17](#)
 - defined [22-3](#)
 - limiting source traffic to specific VLANs [22-22](#)
 - specifying monitored ports [22-17](#)
 - with ingress traffic enabled [22-20](#)
- source ports [22-5](#)
- transmitted traffic [22-5](#)
- VLAN-based [22-6](#)
- RSTP
 - active topology, determining [15-6](#)
 - BPDU
 - format [15-9](#)
 - processing [15-10](#)
 - designated port, defined [15-6](#)
 - designated switch, defined [15-6](#)
 - interoperability with 802.1D
 - described [15-5](#)
 - restarting migration process [15-22](#)
 - topology changes [15-10](#)
 - overview [15-6](#)
 - port roles
 - described [15-6](#)
 - synchronized [15-8](#)
 - proposal-agreement handshake process [15-7](#)
 - rapid convergence
 - described [15-7](#)
 - edge ports and Port Fast [15-7](#)
 - point-to-point links [15-7, 15-22](#)
 - root ports [15-7](#)
 - root port, defined [15-6](#)
 - See also MSTP
- running configuration, saving [4-11](#)

S

- SC (standby command switch) [5-10, 5-18](#)
- scheduled reloads [4-17](#)
- secure MAC addresses
 - deleting [19-12](#)
 - maximum number of [19-8](#)
 - types of [19-7](#)
- secure ports, configuring [19-7](#)
- secure remote connections [7-38](#)
- Secure Shell
 - See SSH
- security, port [19-7](#)
- security features [1-5](#)
- sequence numbers in log messages [24-7](#)
- server mode, VTP [12-3](#)
- service-provider network
 - MSTP and RSTP [15-1](#)
- set-request operation [25-5](#)
- setup program, failed command switch replacement [29-8, 29-10](#)
- severity levels, defining in system messages [24-8](#)
- SFPs
 - security and identification [29-12](#)
- shaped round robin
 - See SRR
- show access-lists hw-summary command [26-17](#)
- show and more command output, filtering [2-9](#)
- show cdp traffic command [20-5](#)
- show cluster members command [5-21](#)
- show configuration command [9-14](#)
- show forward command [29-19](#)
- show interfaces command [9-11, 9-14](#)
- show platform forward command [29-19](#)
- show running-config command
 - displaying ACLs [26-16, 26-17, 26-23, 26-26](#)
 - interface description in [9-14](#)

shutdown command on interfaces [9-17](#)

Simple Network Management Protocol

See SNMP

SmartPort macros

configuration guidelines [10-2](#)

creating and applying [10-3](#)

default configuration [10-2](#)

defined [10-1](#)

displaying [10-4](#)

tracing [10-2](#)

SNAP [20-1](#)

SNMP

accessing MIB variables with [25-5](#)

agent

described [25-4](#)

disabling [25-8](#)

authentication level [25-10](#)

community strings

configuring [25-8](#)

for cluster switches [25-4](#)

overview [25-4](#)

configuration examples [25-15](#)

default configuration [25-7](#)

engine ID [25-7](#)

groups [25-7, 25-9](#)

host [25-7](#)

ifIndex values [25-6](#)

in-band management [1-4](#)

in clusters [5-14](#)

informs

and trap keyword [25-11](#)

described [25-5](#)

differences from traps [25-5](#)

enabling [25-14](#)

limiting access by TFTP servers [25-15](#)

limiting system log messages to NMS [24-9](#)

manager functions [1-3, 25-3](#)

managing clusters with [5-21](#)

MIBs

location of [A-3](#)

supported [A-1](#)

notifications [25-5](#)

overview [25-1, 25-5](#)

security levels [25-3](#)

status, displaying [25-16](#)

system contact and location [25-14](#)

trap manager, configuring [25-13](#)

traps

described [25-3, 25-5](#)

differences from informs [25-5](#)

enabling [25-11](#)

enabling MAC address notification [6-24](#)

overview [25-1, 25-5](#)

types of [25-11](#)

users [25-7, 25-9](#)

versions supported [25-2](#)

SNMPv1 [25-2](#)

SNMPv2C [25-2](#)

SNMPv3 [25-2](#)

snooping, IGMP [18-2](#)

software images

location in Flash [B-20](#)

recovery procedures [29-2](#)

scheduling reloads [4-17](#)

tar file format, described [B-21](#)

See also downloading and uploading

source addresses, in ACLs [26-9](#)

source-and-destination-IP address based forwarding,
EtherChannel [28-7](#)

source-and-destination MAC address forwarding,
EtherChannel [28-7](#)

source-IP address based forwarding, EtherChannel [28-7](#)

source-MAC address forwarding, EtherChannel [28-7](#)

SPAN

configuration guidelines [22-10](#)

default configuration [22-9](#)

- destination ports [22-7](#)
- displaying status [22-23](#)
- interaction with other features [22-8](#)
- monitored ports [22-5](#)
- monitoring ports [22-7](#)
- overview [1-7](#), [22-1](#)
- received traffic [22-4](#)
- session limits [22-10](#)
- sessions
 - configuring ingress forwarding [22-14](#), [22-21](#)
 - creating [22-11](#)
 - defined [22-3](#)
 - limiting source traffic to specific VLANs [22-15](#)
 - removing destination (monitoring) ports [22-12](#)
 - specifying monitored ports [22-11](#)
 - with ingress traffic enabled [22-13](#)
- source ports [22-5](#)
- transmitted traffic [22-5](#)
- VLAN-based [22-6](#)
- spanning tree and native VLANs [11-17](#)
- Spanning Tree Protocol
 - See STP
- SPAN traffic [22-4](#)
- speed, configuring on interfaces [9-10](#)
- SRR
 - configuring
 - shaped weights on egress queues [27-61](#)
 - shared weights on egress queues [27-62](#)
 - shared weights on ingress queues [27-55](#)
 - described [27-12](#)
 - shaped mode [27-12](#)
 - shared mode [27-12](#)
 - support for [1-7](#)
- SSH
 - configuring [7-39](#)
 - cryptographic software image [7-37](#)
 - described [1-4](#), [7-38](#)
 - encryption methods [7-38](#)
 - user authentication methods, supported [7-38](#)
- Standby Command Configuration window [5-19](#)
- standby command switch
 - configuring [5-18](#)
 - considerations [5-11](#)
 - defined [5-2](#)
 - priority [5-10](#)
 - requirements [5-3](#)
 - virtual IP address [5-11](#)
- See also cluster standby group and HSRP
- standby group, cluster
 - See cluster standby group and HSRP
- startup configuration
 - booting
 - manually [4-13](#)
 - specific image [4-14](#)
 - clearing [B-19](#)
 - configuration file
 - automatically downloading [4-13](#)
 - specifying the filename [4-13](#)
 - default boot configuration [4-12](#)
- static access ports
 - assigning to VLAN [11-11](#)
 - defined [9-3](#), [11-3](#)
- static addresses
 - See addresses
- static MAC addressing [1-6](#)
- static VLAN membership [11-2](#)
- statistics
 - 802.1X [8-19](#)
 - CDP [20-5](#)
 - interface [9-16](#)
 - QoS ingress and egress [27-65](#)
 - RMON group Ethernet [23-6](#)
 - RMON group history [23-5](#)
 - SNMP input and output [25-16](#)
 - VTP [12-15](#)

- sticky learning [19-8](#)
- storm control
 - configuring [19-3](#)
 - described [19-2](#)
 - displaying [19-15](#)
 - support for [1-3](#)
 - thresholds [19-2](#)
- STP
 - 802.1D and bridge ID [14-4](#)
 - 802.1D and multicast addresses [14-8](#)
 - 802.1T and VLAN identifier [14-4](#)
 - accelerating root port selection [16-4](#)
 - BackboneFast
 - described [16-5](#)
 - enabling [16-13](#)
 - BPDU filtering
 - described [16-3](#)
 - enabling [16-12](#)
 - BPDU guard
 - described [16-3](#)
 - enabling [16-11](#)
 - BPDU message exchange [14-3](#)
 - configuration guidelines [14-12, 16-9](#)
 - configuring
 - forward-delay time [14-21](#)
 - hello time [14-20](#)
 - maximum aging time [14-21](#)
 - path cost [14-18](#)
 - port priority [14-17](#)
 - root switch [14-14](#)
 - secondary root switch [14-16](#)
 - spanning-tree mode [14-13](#)
 - switch priority [14-19](#)
 - counters, clearing [14-22](#)
 - default configuration [14-11](#)
 - default optional feature configuration [16-9](#)
 - designated port, defined [14-3](#)
 - designated switch, defined [14-3](#)
 - detecting indirect link failures [16-6](#)
 - disabling [14-14](#)
 - displaying status [14-22](#)
 - extended system ID
 - effects on root switch [14-14](#)
 - effects on the secondary root switch [14-16](#)
 - overview [14-4](#)
 - unexpected behavior [14-14](#)
 - features supported [1-4](#)
 - inferior BPDU [14-3](#)
 - instances supported [14-9](#)
 - interface state, blocking to forwarding [16-2](#)
 - interface states
 - blocking [14-5](#)
 - disabled [14-7](#)
 - forwarding [14-5, 14-6](#)
 - learning [14-6](#)
 - listening [14-6](#)
 - overview [14-4](#)
 - interoperability and compatibility among modes [14-10](#)
 - limitations with 802.1Q trunks [14-10](#)
 - load sharing
 - overview [11-22](#)
 - using path costs [11-24](#)
 - using port priorities [11-22](#)
 - loop guard
 - described [16-8](#)
 - enabling [16-15](#)
 - modes supported [14-9](#)
 - multicast addresses, effect of [14-8](#)
 - optional features supported [1-5](#)
 - overview [14-2](#)
 - path costs [11-24, 11-25](#)
 - Port Fast
 - described [16-2](#)
 - enabling [16-10](#)
 - port priorities [11-23](#)
 - preventing root switch selection [16-7](#)
 - protocols supported [14-9](#)
 - redundant connectivity [14-8](#)

- root guard
 - described 16-7
 - enabling 16-14
- root port, defined 14-3
- root switch
 - configuring 14-14
 - effects of extended system ID 14-4, 14-14
 - election 14-3
 - unexpected behavior 14-14
- shutdown Port Fast-enabled port 16-3
- status, displaying 14-22
- superior BPDU 14-3
- timers, described 14-20
- UplinkFast
 - described 16-4
 - enabling 16-13
- stratum, NTP 6-2
- success response, VMPS 11-26
- summer time 6-14
- SunNet Manager 1-3
- switch clustering technology 5-1
 - See also clusters, switch 1-2
 - See clusters, switch
- switch console port 1-4
- Switched Port Analyzer
 - See SPAN
- switched ports 9-2
- Switch Manager 3-14
 - See also Device Manager
- switchport block multicast command 19-6
- switchport block unicast command 19-6
- switchport protected command 19-5
- switch priority
 - MSTP 15-19
 - STP 14-19
- switch software features 1-1
- syslog
 - See system message logging
- system clock
 - configuring
 - daylight saving time 6-14
 - manually 6-12
 - summer time 6-14
 - time zones 6-13
 - displaying the time and date 6-12
 - overview 6-2
 - See also NTP
- system message logging
 - default configuration 24-3
 - defining error message severity levels 24-8
 - disabling 24-4
 - displaying the configuration 24-12
 - enabling 24-4
 - facility keywords, described 24-12
 - level keywords, described 24-8
 - limiting messages 24-9
 - message format 24-2
 - overview 24-1
 - sequence numbers, enabling and disabling 24-7
 - setting the display destination device 24-4
 - synchronizing log messages 24-5
 - syslog facility 1-7
 - time stamps, enabling and disabling 24-7
 - UNIX syslog servers
 - configuring the daemon 24-10
 - configuring the logging facility 24-11
 - facilities supported 24-12
- system name
 - default configuration 6-16
 - default setting 6-16
 - manual configuration 6-16
 - See also DNS
- system prompt
 - default setting 6-16
 - manual configuration 6-17

T

TACACS+

- accounting, defined [7-11](#)
- authentication, defined [7-11](#)
- authorization, defined [7-11](#)
- configuring
 - accounting [7-17](#)
 - authentication key [7-13](#)
 - authorization [7-16](#)
 - login authentication [7-14](#)
- default configuration [7-13](#)
- displaying the configuration [7-17](#)
- identifying the server [7-13](#)
- in clusters [5-14](#)
- limiting the services to the user [7-16](#)
- operation of [7-12](#)
- overview [7-10](#)
- support for [1-6](#)
- tracking services accessed by user [7-17](#)

tar files

- creating [B-6](#)
- displaying the contents of [B-6](#)
- extracting [B-7](#)
- image file format [B-21](#)

TDR [1-7](#)

Telnet

- accessing management interfaces [2-10](#)
- from a browser [2-10](#)
- number of connections [1-4](#)
- setting a password [7-6](#)

Terminal Access Controller Access Control System Plus

See TACACS+

terminal lines, setting a password [7-6](#)

TFTP

- configuration files
 - downloading [B-11](#)

preparing the server [B-10](#)

uploading [B-11](#)

configuration files in base directory [4-6](#)

configuring for autoconfiguration [4-6](#)

image files

- deleting [B-24](#)
- downloading [B-23](#)
- preparing the server [B-22](#)
- uploading [B-24](#)

limiting access by servers [25-15](#)

TFTP server [1-4](#)

threshold, traffic level [19-2](#)

time

See NTP and system clock

Time Domain Reflector

See TDR

time-range command [26-13](#)

time ranges in ACLs [26-13](#)

time stamps in log messages [24-7](#)

time zones [6-13](#)

Token Ring VLANs

- support for [11-5](#)
- VTP support [12-4](#)

Topology view

described [3-2, 3-14](#)

TOS [1-6](#)

traceroute, Layer 2

- and ARP [29-15](#)
- and CDP [29-14](#)
- described [29-14](#)
- IP addresses and subnets [29-15](#)
- MAC addresses and VLANs [29-15](#)
- multicast traffic [29-15](#)
- multiple devices on a port [29-15](#)
- unicast traffic [29-14](#)
- usage guidelines [29-14](#)

traceroute command [29-16](#)

See also IP traceroute

- traffic
 - blocking flooded [19-6](#)
 - fragmented [26-4](#)
 - unfragmented [26-4](#)
 - traffic policing [1-7](#)
 - traffic suppression [19-2](#)
 - transparent mode, VTP [12-3, 12-11](#)
 - trap-door mechanism [4-2](#)
 - traps
 - configuring MAC address notification [6-24](#)
 - configuring managers [25-11](#)
 - defined [25-3](#)
 - enabling [6-24, 25-11](#)
 - notification types [25-11](#)
 - overview [25-1, 25-5](#)
 - troubleshooting
 - connectivity problems [29-13, 29-14, 29-16](#)
 - detecting unidirectional links [21-1](#)
 - determining packet forwarding [29-19](#)
 - displaying crash information [29-21](#)
 - SFP security and identification [29-12](#)
 - show forward command [29-19](#)
 - with CiscoWorks [25-5](#)
 - with debug commands [29-17](#)
 - with ping [29-13](#)
 - with system message logging [24-1](#)
 - with traceroute [29-16](#)
 - trunking encapsulation [1-5](#)
 - trunk ports
 - configuring [11-19](#)
 - defined [9-3, 11-3](#)
 - encapsulation [11-19, 11-23, 11-25](#)
 - secure MAC addresses on [19-11](#)
 - trunks
 - allowed-VLAN list [11-20](#)
 - configuring [11-19, 11-23, 11-25](#)
 - ISL [11-15](#)
 - load sharing
 - setting STP path costs [11-24](#)
 - using STP port priorities [11-22, 11-23](#)
 - native VLAN for untagged traffic [11-21](#)
 - parallel [11-24](#)
 - pruning-eligible list [11-21](#)
 - to non-DTP device [11-16](#)
 - understanding [11-15](#)
 - trusted boundary for QoS [27-34](#)
 - trusted port states
 - between QoS domains [27-35](#)
 - classification options [27-5](#)
 - ensuring port security for IP phones [27-34](#)
 - support for [1-6](#)
 - within a QoS domain [27-31](#)
 - twisted-pair Ethernet, detecting unidirectional links [21-1](#)
 - type of service
 - See TOS
-
- ## U
- UDLD
 - default configuration [21-4](#)
 - echoing detection mechanism [21-3](#)
 - enabling
 - globally [21-5](#)
 - per interface [21-5](#)
 - link-detection mechanism [21-1](#)
 - neighbor database [21-2](#)
 - overview [21-1](#)
 - resetting an interface [21-6](#)
 - status, displaying [21-6](#)
 - support for [1-4](#)
 - unauthorized ports with 802.1X [8-4](#)
 - unicast MAC address filtering [1-4](#)
 - and adding static addresses [6-28](#)

- and broadcast MAC addresses [6-27](#)
- and CPU packets [6-27](#)
- and multicast addresses [6-27](#)
- and router MAC addresses [6-27](#)
- configuration guidelines [6-27](#)
- described [6-27](#)
- unicast storm control command [19-4](#)
- unicast storms [19-2](#)
- unicast traffic, blocking [19-6](#)
- UniDirectional Link Detection protocol
 - See UDLD
- UNIX syslog servers
 - daemon configuration [24-10](#)
 - facilities supported [24-12](#)
 - message logging configuration [24-11](#)
- unrecognized Type-Length-Value (TLV) support [12-4](#)
- upgrading software images
 - See downloading
- UplinkFast
 - described [16-4](#)
 - enabling [16-13](#)
 - support for [1-5](#)
- uploading
 - configuration files
 - preparing [B-10](#), [B-13](#), [B-16](#)
 - reasons for [B-8](#)
 - using FTP [B-15](#)
 - using RCP [B-18](#)
 - using TFTP [B-11](#)
 - image files
 - preparing [B-22](#), [B-25](#), [B-29](#)
 - reasons for [B-20](#)
 - using FTP [B-28](#)
 - using RCP [B-33](#)
 - using TFTP [B-24](#)
- user EXEC mode [2-2](#)
- username-based authentication [7-7](#)

V

- version-dependent transparent mode [12-4](#)
- virtual IP address
 - cluster standby group [5-11](#), [5-19](#)
 - command switch [5-11](#), [5-19](#)
 - See also IP addresses
- vlan.dat file [11-4](#)
- VLAN 1, disabling on a trunk port [11-20](#)
- VLAN 1 minimization [11-20](#)
- VLAN ACLs
 - See VLAN maps
- vlan-assignment response, VMPS [11-26](#)
- VLAN configuration
 - at bootup [11-7](#)
 - saving [11-7](#)
- VLAN configuration mode [2-2](#), [11-7](#)
- VLAN database
 - and startup configuration file [11-7](#)
 - and VTP [12-1](#)
 - VLAN configuration saved in [11-7](#)
 - VLANs saved in [11-4](#)
- vlan database command [11-7](#)
- VLAN filtering, and SPAN [22-6](#)
- vlan global configuration command [11-6](#)
- VLAN ID, discovering [6-29](#)
- VLAN management domain [12-2](#)
- VLAN Management Policy Server
 - See VMPS
- VLAN map entries, order of [26-23](#)
- VLAN maps
 - applying [26-26](#)
 - common uses for [26-26](#)
 - configuration example [26-27](#)
 - configuration guidelines [26-23](#)
 - configuring [26-22](#)
 - creating [26-23](#)

- defined 26-2, 26-3
- denying access example 26-28
- denying and permitting packets 26-24
- displaying 26-29
- examples 26-28
- support for 1-6
- with router ACLs 26-29
- VLAN membership
 - confirming 11-29
 - modes 11-3
- VLAN Query Protocol
 - See VQP
- VLANs
 - adding 11-8
 - adding to VLAN database 11-8
 - aging dynamic addresses 14-9
 - allowed on trunk 11-20
 - and spanning-tree instances 11-3, 11-6, 11-13
 - configuration guidelines, extended-range VLANs 11-12
 - configuration guidelines, normal-range VLANs 11-6
 - configuration options 11-6
 - configuring 11-1
 - configuring IDs 1006 to 4094 11-12
 - creating in config-vlan mode 11-9
 - creating in VLAN configuration mode 11-10
 - default configuration 11-7
 - deleting 11-10
 - described 9-2, 11-1
 - displaying 11-14
 - extended-range 11-1, 11-12
 - features 1-5
 - illustrated 11-2
 - limiting source traffic with RSPAN 22-22
 - limiting source traffic with SPAN 22-15
 - modifying 11-8
 - native, configuring 11-21
 - normal-range 11-1, 11-4
 - number supported 1-5
 - parameters 11-5
 - port membership modes 11-3
 - static-access ports 11-11
 - STP and 802.1Q trunks 14-10
 - supported 11-3
 - Token Ring 11-5
 - traffic between 11-2
 - VTP modes 12-3
- VLAN Trunking Protocol
 - See VTP
- VLAN trunks 11-15
- VMPS
 - administering 11-30
 - configuration example 11-31
 - configuration guidelines 11-27
 - default configuration 11-27
 - description 11-26
 - dynamic port membership
 - described 11-27
 - reconfirming 11-29
 - troubleshooting 11-31
 - entering server address 11-28
 - mapping MAC addresses to VLANs 11-26
 - monitoring 11-30
 - reconfirmation interval, changing 11-29
 - reconfirming membership 11-29
 - retry count, changing 11-30
- voice-over-IP 13-1
- voice VLAN
 - Cisco 7960 phone, port connections 13-1
 - configuration guidelines 13-3
 - configuring IP phones for data traffic
 - override CoS of incoming frame 13-5
 - trust CoS priority of incoming frame 13-5
 - configuring ports for voice traffic in
 - 802.1P priority tagged frames 13-5
 - 802.1Q frames 13-4

- connecting to an IP phone [13-4](#)
- default configuration [13-3](#)
- described [13-1](#)
- displaying [13-6](#)
- VQP [1-5, 11-26](#)
- VTP
 - adding a client to a domain [12-14](#)
 - advertisements [11-18, 12-3](#)
 - and extended-range VLANs [12-1](#)
 - and normal-range VLANs [12-2](#)
 - client mode, configuring [12-10](#)
 - configuration
 - global configuration mode [12-7](#)
 - guidelines [12-7](#)
 - privileged EXEC mode [12-7](#)
 - requirements [12-8](#)
 - saving [12-7](#)
 - VLAN configuration mode [12-7](#)
 - configuration mode options [12-6](#)
 - configuration requirements [12-8](#)
 - configuration revision number
 - guideline [12-14](#)
 - resetting [12-14](#)
 - configuring
 - client mode [12-10](#)
 - server mode [12-9](#)
 - transparent mode [12-11](#)
 - consistency checks [12-4](#)
 - default configuration [12-6](#)
 - described [12-1](#)
 - disabling [12-11](#)
 - domain names [12-7](#)
 - domains [12-2](#)
 - modes
 - client [12-3, 12-10](#)
 - server [12-3, 12-9](#)
 - transitions [12-3](#)
 - transparent [12-3, 12-11](#)
 - monitoring [12-15](#)

- passwords [12-8](#)
- pruning
 - disabling [12-13](#)
 - enabling [12-13](#)
 - examples [12-5](#)
 - overview [12-4](#)
 - support for [1-5](#)
- pruning-eligible list, changing [11-21](#)
- server mode, configuring [12-9](#)
- statistics [12-15](#)
- support for [1-5](#)
- Token Ring support [12-4](#)
- transparent mode, configuring [12-11](#)
- using [12-1](#)
- version, guidelines [12-8](#)
- version 1 [12-4](#)
- version 2
 - configuration guidelines [12-8](#)
 - disabling [12-13](#)
 - enabling [12-12](#)
 - overview [12-4](#)

W

- weighted tail drop
 - See WTD
- wizards [1-2, 3-6](#)
- WTD
 - described [27-11](#)
 - setting thresholds
 - egress queue-sets [27-58](#)
 - ingress queues [27-53](#)
 - support for [1-7](#)

X

- XMODEM protocol [29-2](#)